A Fundamental Scalability Criterion for Data Aggregation in VANETs

Björn Scheuermann Christian Lochert Jedrzej Rybicki Martin Mauve
Computer Networks Group
Heinrich Heine University
Düsseldorf, Germany
{scheuermann, lochert, rybicki, mauve}@cs.uni-duesseldorf.de

ABSTRACT

The distribution of dynamic information from many sources to many destinations is a key challenge for VANET applications such as cooperative traffic information management or decentralized parking guidance systems. In order for these systems to remain scalable it has been proposed to aggregate the information within the network as it travels from the sources to the destinations. However, so far it has remained unclear by what amount the aggregation scheme needs to reduce the original data in order to be considered scalable. In this paper we prove formally that any suitable aggregation scheme must reduce the bandwidth at which information about an area at distance d is provided to the cars asymptotically faster than $1/d^2$. Furthermore, we constructively show that this bound is tight: for any arbitrary $\epsilon > 0$, there exists a scalable aggregation scheme that reduces information asymptotically like $1/d^{2+\epsilon}$.

Categories and Subject Descriptors

C.4 [**Computer Systems Organization**]: Performance of Systems; C.2.1 [**Computer Systems Organization**]: Computer-Communication Networks—*Network Architecture and Design*; H.4.3 [**Information Systems**]: Information Systems Applications—*Communications Applications*

General Terms

Algorithms, Design, Theory

Keywords

VANET, Aggregation, Data Dissemination, Scalability, Capacity, Bandwidth Profile

1. INTRODUCTION

Research on Vehicular Ad-hoc Networks (VANETs) aims to improve road safety, traffic efficiency, and driver convenience. Well-known examples in the area of efficiency and convenience applications are cooperative traffic information management and parking guidance systems. This class of applications typically requires the distribution of dynamic information (such as traffic condition or free parking slots) originating from many locations to many or all nodes in a large network area. One key communication paradigm in VANETs is therefore the continuous transmission of measurement data from many sources to many destinations.

Yet, VANETs, just like wireless multihop networks in general, have very limited capacity. Hence, it is obviously not possible to send continuous updates about each location where measurements are taken to all network participants at a fixed data rate. It has thus been proposed to aggregate data more and more with increasing distance, i. e., to maintain and distribute a detailed picture within the closer vicinity and coarser and coarser information about increasingly distant regions.

However, while many aggregation mechanisms and applications based on aggregation have been proposed—examples are [3, 13–15]—little is understood so far about the fundamental limitations and requirements of VANET data aggregation. It has often been stated that aggregation is necessary for scalable VANET information dissemination. But what are the characteristics of suitable aggregation schemes? How frequently can updates of, for example, traffic information be provided to remote cars? By how much do we have to reduce the "resolution" of information about distant regions? I. e., how much "coarser" does the picture have to be and how much aggregation is "enough"?

This paper is meant as a step towards a deeper understanding of these fundamental issues. Previous work has considered individual, specific aggregation schemes and dissemination mechanisms, and has evaluated them in specific situations and environments, using simulations and experiments. Here, we consider the generic class of *all* possible aggregation and dissemination mechanisms. We are interested in capturing the underlying effects and principles, in order to derive general limits that *any* protocol design must respect. Consequently, our methodology of choice is not simulation or experimentation, but theoretical analysis.

To deal with the broad range of conceivable aggregation schemes, we need a suitable abstraction that captures the essence of in-network data aggregation regardless how it is specifically performed. To this end, we introduce the notion of *bandwidth profiles*. Speaking simplified (a rigorous definition follows), a bandwidth profile of an aggregation scheme describes how rapidly the amount of information made available is reduced with increasing distance. As it turns out, bandwidth profiles constitute a valuable tool to describe the general properties of aggregation schemes; all our main results will be formulated in terms of bandwidth profiles.

Our primary focus is on the minimum aggregation requirements for scalable dissemination applications in two-dimensional wireless networks—a setting which describes VANET dissemination applications well. We consider this from an asymptotic perspective, demonstrating that a network and application model with very weak assumptions already allows to derive interesting results. Because our assumptions on the network are weak, the results are strong: the proofs hold for a broad class of protocols and aggregation algorithms.

In particular, we show that any dissemination mechanism, in order to be scalable in a general setting, must reduce the bandwidth at which information about an area at distance d is provided to the cars asymptotically faster than $1/d^2$. This result does not depend on the *way* how this bandwidth reduction is achieved: it holds regardless of whether, for instance, information about distant roads, parking lots, etc. is updated at a lower frequency, whether data from multiple sources is summarized, whether less accurate (and thus more compact) data representations are employed, or whether some or all of these techniques are combined in order to reduce the utilized bandwidth. It also does not depend on the communication paradigm used for transporting the information, be it proactive or request-based, using unicast, multicast, broadcast, geocast, DTN-style opportunistic gossiping, or anything else. We subsequently argue that aggregation mechanisms can actually come arbitrarily close to the above stated limit of $o(1/d^2)$: we show how aggregation schemes could assign the available bandwidth to information sources in order to achieve an asymptotic behavior like $1/d^{2+\epsilon}$ for any arbitrary $\epsilon > 0$.

Our results provide valuable hints on how specific aggregation mechanisms should be designed. To name just a few exemplary implications: The self-organizing traffic information system (SOTIS) [24] summarizes information about preconfigured, fixed-size road segments. CASCADE [8] uses syntactic compression to reduce the size of individual records. Both mechanisms can be used as components in an overall system design. However, in the light of the asymptotic data reduction properties which we prove to be necessary for general scalability, it becomes clear that summaries of fixed-size road segments or a size reduction of individual records alone do not suffice. In TrafficView [15], more and more traffic information records are merged with increasing distance. Two algorithms for this merging process are proposed, both can be configured by a number of parameters. Our result provides hints on how to choose these parameters: the size of the aggregated records should be reduced more quickly than $1/d^2$ over distance. In [13], we discuss how an area-based aggregation scheme can be built, but the specific structure of the regions remains open. Our work here provides a guideline on how quickly the size of these regions should increase over distance. The results of our analysis may also be applicable outside the VANET context. One example where this is particularly easy to see is the location dissemination mechanism in the DREAM routing protocol [2]. In DREAM, the node positions are announced with varying frequency and range. While DREAM does not use any in-network aggregation, its use of network resources for position announcements can still be characterized by means of bandwidth profiles; these must obey the limitations proven here in order to scale well with increasing network size.

The remainder of this paper is structured as follows. In Section 2 we introduce a model of data dissemination as a basis for further considerations. In Section 3 we then prove the abovementioned bound on the necessary bandwidth reduction. That it can be approached arbitrarily closely is shown in Section 4. Subsequently, we discuss implications and applicability of our results in Section 5 and review related work in Section 6. Finally, we conclude with a summary in Section 7.

2. MODEL

In this section, we introduce the network and application model that we will use throughout the paper. Our aim is to capture the relevant aspects of wireless networks in general and VANETs in particular, while keeping the focus of our assumptions on those factors that are later on essential for our proofs—each non-essential constraint in the model would unnecessarily limit the applicability and generality of the results. Specifically, our model comprises three components:

1. The sources of information, i. e., where the disseminated and aggregated data comes from. Here, we call these sources *measurement points*.

2. Where the information goes, i. e., which information is to be delivered to which cars in the VANET. We term this relation the *interests* in the system.

3. The limitations on the propagation of information imposed by the network as a result of limited spatial reuse of the medium, in particular due to wireless interference.

2.1 Measurement points

Our model represents the "world", i. e., the area the system is deployed on, by the real plane $\mathbb{R}^2$. For practical purposes, this is a reasonable approximation of a city area, a country, or even a continent. On this plane, there is a set $M \subset \mathbb{R}^2$ of locations at which information can be obtained through measurements. These *measurement points* could, for example, represent all the street segments, for which passing-by cars would observe the current traffic density, driving velocity, number of free parking places, road surface condition, or any other parameter. The observed values are time-varying, i. e., the measurements are always taken at some specific time instant. Due to this temporal property, we see a measurement point as an information source that "produces" information about the measured value whenever a measurement is performed. Considered over a sufficiently long time span and seen from an abstract perspective, a measurement point thus is a source which generates information at a certain *data rate*. The task of an information dissemination protocol is to deliver this generated information to the interested network participants. The focus here is to assess the asymptotic limits of the rate at which information obtained from the measurement points can be *delivered* to interested cars, regardless of the specific protocols used to transport it and the in-network aggregation techniques used to achieve the desired rate.

While, at a first glance, this abstraction bears similarities with existing work on asymptotic rate limits, there is one fundamentally distinctive aspect: with in-network aggregation, the data rate changes *within* the network. Typically, with increasing distance, the amount of information about a measurement point is reduced further and further, for instance by summarizing it with other data.

Generally, the achievable performance will depend on the distribution and density of the measurement points on the plane. If there are only few measurement points, disseminating the information will be easier than in the case of a large number of information sources. We therefore have to model the spatial distribution of the measurement points. If we allow for arbitrarily large and dense clusters of measurement points, an arbitrarily large amount of information can be generated in a very limited area; then, their information can obviously not be communicated even locally and asymptotic bandwidth considerations become meaningless. However, since we intend to stay as generic as possible in our analysis, we impose only a very weak condition on the distribution of the measurement points. We concentrate on the case where the measurement point distribution satisfies what we call a *max-density condition.* This condition essentially states that there are no arbitrarily large, arbitrarily dense groups or clusters of measurement points. It is formally defined as follows:

DEFINITION 1. *A set of measurement points M fulfills a max-density condition with parameters $\delta > 0$ and $r_0 > 0$ if and only if for any circle in $\mathbb{R}^2$ with radius $r \geq r_0$ the number ν of measurement points that lie within the circle is bounded above by*

$$\nu < \delta r^2.$$

Note that in the previous definition, parameter choices where $\delta r_0^2 \leq 1$ do not make sense, because they would not allow for even just one single measurement point to exist: let m be a measurement point and consider a circle with radius $r = r_0$ around m; this circle would contain one measurement point and would thus already violate a max-density condition with $\delta r_0^2 \leq 1$. Therefore, we may safely assume that δr_0^2 will always be larger than 1.

2.2 Interests

The next aspect that needs to be modeled is the distribution of interests in the system: where does information need to be delivered? Or, more specifically, cars in which region of the map are interested in information about which measurement points? This aspect is highly application dependent. In a traffic information system, for example, cars will be interested in the traffic situation along the whole of their planned route; thus, the interests depend on the traffic movement pattern, and will typically include both close-by measurement points and much more distant road sections. Because our aim is to investigate the fundamental limits that hold for any application and any road network, our model does intentionally not constrain the possible distribution of interests in the system. We therefore model the interests of the participants in the dissemination system by an arbitrary set of pairs (x, m), where x is the location of a network participant that is interested in data from a measurement point $m \in M$. The formal definition is as follows.

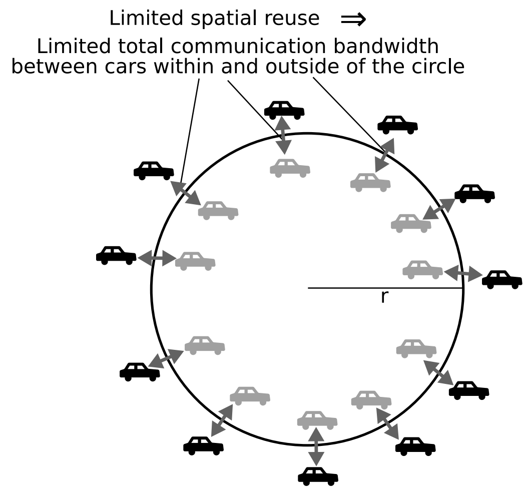

Figure 1: The total communication bandwidth into a circle of finite radius is finite.

DEFINITION 2. *Let the interest set $\mathcal{I}$ be an arbitrary subset of $\mathbb{R}^2 \times M$.*

Since we are interested in the rate at which data about a measurement point can be made available given the distance between the measurement point and the interested car, we also introduce a notion for this distance:

DEFINITION 3. *The* distance $\|I\|$ *of an interest $I \in \mathcal{I}$, where $I = (x, m)$, is the distance between the position of the interested party x and the measurement point m, i. e.,*

$$\|I\| := \|x - m\|.$$

2.3 Limitations of the network

Finally, we have to consider the network capacity itself. In a network with unlimited capacity, it would not be a problem to deliver all measurement data about all measurement points to each interested participant. In practice, however, each network imposes limitations on the maximum bandwidth between communication partners. In wireless networks like VANETs, the central limiting factor are spatial reuse constraints due to signal interference. In order to obtain strong results, we again aim to capture the essence of these limitations in a way that is as generic as possible, with as few specifics of and assumptions about any particular network or communication mechanism as possible.

We again do so using circles on the plane. Any other shape would do equally well, but circles are particularly easy to handle. The idea is sketched in Figure 1: we place a circle with finite radius somewhere in our wireless network. If we add up all the bandwidth of all communication links crossing the circle boundary that we could possibly use in parallel, this sum will always be finite—wireless interference limits the spatial reuse. We thus cannot communicate an arbitrarily large amount of data into (or out of) our circle within limited time.

We formalize this observation in the following assumption. Note that it only constitutes an upper bound; we do not assume that the maximum bandwidth is ever actually achieved.

Assumption 1. *Given any arbitrary radius $r > 0$, we assume that there is a constant ξ_r, such that for any circle with radius r, the total communication bandwidth b (observed over a reasonably long timespan) between the nodes within and outside the circle is bounded above by*

$$b \leq \xi_r.$$

The total communication bandwidth between the inside and the outside of our circle (i. e., the constant ξ_r) will typically increase for larger radii r. However, we do not need to model this in detail: for our proof, it is sufficient that for any fixed, given radius, the total bandwidth may not become arbitrarily high.

3. BANDWIDTH CONSTRAINTS FOR SCALABLE AGGREGATION

Before we turn towards the question of *how much* aggregation is necessary, let us first argue—in terms of our model—why data aggregation is necessarily needed for dissemination services in VANETs *at all*. Consider a circle with arbitrary radius r around the participant's current position; according to Assumption 1 the total data rate of information from measurement points outside this circle provided to the participant cannot exceed ξ_r. If each measurement point produces data with at least some positive minimum data rate and the participant is interested in more and more measurement points outside the circle, the necessary total data rate *will* at some point exceed ξ_r. Thus, in order to be able to serve the interests of the network participants under all circumstances, a VANET dissemination protocol *must* use techniques to reduce the data rate at which information about distant regions is provided.

One approach would of course be to limit the number of interests. This, however, is hardly viable. It is of course easily possible to put a limit on the allowed number of interests of a single network participant; such a limit may even inherently exist in the application. Note, however, that it will often be the case that many network participants which are interested in very different regions are co-located within the same geographical area. For example, cars with very different destinations and different planned routes may be underway on the same road. However, despite the almost arbitrarily large variety of interests, the total ingress bandwidth for these cars is still limited. (In terms of our model, consider a circle enclosing all the cars, then the total bandwidth into this circle is limited.) Thus, a much more promising approach is to use in-network summarization and aggregation techniques to adjust the resolution of the provided information (and thus the necessary data rate) in order to respect the inherent bandwidth limitations of inter-vehicle communication.

3.1 Bandwidth profiles

Aggregation techniques exist in many different shades and flavors. For instance, information from the measurement points within the same geographical area may be aggregated, such that only a summary is transmitted to interested parties further away, as it has been proposed in [3,13]; information from individual vehicles on a road may be combined as in [15]; the frequency at which information about certain geographical regions is transmitted may be adjusted like in [11]; syntactic compression techniques may be used to yield smaller aggregates, as suggested in [8]; or travel times between important "landmarks" may be used as a coarser representation of the traffic situation at larger distances as it has been proposed in [14].

All these techniques—more or less directly—aim at reducing the data rate used per measurement point with increasing distance of the interest, either by using coarser approximations of the value itself (i. e., data representations requiring less bits), by sending updates less often, or by transmitting one single value for a whole group of measurement points, which in the end also boils down to reducing the rate used per measurement point.

In our analytical approach to determine the fundamental limits of aggregation, we must find a way to abstract from the specific approaches and mechanisms of aggregation. We therefore identify the reduction of the data rate per measurement point with increasing distance as their common feature. From this point of view, the essential property of an aggregation scheme is the amount of bandwidth spent depending on the distance: given an interest with distance d, how much bandwidth is used for making information available to the interested party? Consequently, we characterize an aggregation scheme by its *bandwidth profile*.

Definition 4. *A monotonically decreasing function*

$$b: \mathbb{R} \to \mathbb{R}_{\geq 0}$$

is a bandwidth profile of an aggregation scheme $\mathcal{A}$ if $\mathcal{A}$ ensures that for all interests $I = (x, m) \in \mathcal{I}$ an interested party located at x is supplied with information about measurement point m at least with data rate $b(\|I\|)$.

3.2 Towards a proof

In the remainder of this section, we will prove for the general case of arbitrary measurement point sets (with a max-density property) and interests that if an aggregation scheme has a bandwidth profile b for which $b(d) \notin o(1/d^2)$, then the bandwidth limitations established by Assumption 1 do not allow to ensure that all interests can be appropriately served.[1] I. e., we will prove by contradiction that all bandwidth profiles of scalable aggregation schemes must be in $o(1/d^2)$. Hence, practical aggregation schemes should be constructed in a way such that the data rate used per measurement point decreases faster than with the square of the interest's distance.

In the proof, we will make use of the following two lemmas.

Lemma 1. *Let $\delta > 0$, $r_0 > 0$ be given such that $\delta r_0^2 > 1$. If the measurement points in M are distributed in such a way that the distance between any two measurement points is at least*

$$\Delta := \frac{2}{\sqrt{\delta} - r_0^{-1}},$$

then a max-density condition with parameters δ and r_0 according to Definition 1 holds.

[1] $o(1/d^2)$, not to be confused with $O(1/d^2)$, describes the set of functions which decrease asymptotically faster than $1/d^2$. More formally, $b(d) \in o(1/d^2)$ if and only if

$$\forall c > 0 : \exists d_0 > 0 : \forall d > d_0 : b(d) < \frac{c}{d^2}.$$

The proof of this lemma can be found in Appendix A.

Note that the condition in Lemma 1 is sufficient, but not necessary. In particular, we do not claim that measurement points will always have a pair-wise distance of Δ; we will just use this lemma as a vehicle to show that one specific distribution of measurement points (the counterexample in the proof by contradiction of our main theorem) fulfills a max-density condition.

Lemma 2. *On the perimeter of a circle with radius $r \geq \frac{\Delta}{2}$, at least $\left\lfloor \frac{4r}{\Delta} \right\rfloor$ points can be positioned such that for each pair of points their distance is at least Δ.*

This lemma is proven in Appendix B.

3.3 Constructing a counterexample

In the following, we will assume that the parameters δ and r_0 from the max-density condition are fixed and given. Let furthermore Δ in the following denote the quantity $\frac{2}{\sqrt{\delta - r_0^{-1}}}$, as in Lemma 1.

We will now show that any bandwidth profile must be in $o(1/d^2)$, or otherwise there is a possible constellation of measurement points and interests such that a max-density condition holds, but not all interests I can be satisfied with data rate $b(\|I\|)$. We prove this by contradiction. Therefore, let from now on b be a bandwidth profile for which

$$b(d) \notin o\left(\frac{1}{d^2}\right). \tag{1}$$

We will proceed in three steps towards a proof. We first construct a set of measurement points M^* and set of interests $\mathcal{I}^*$ based on the parameters δ, r_0 from the max-density condition. We then show that for this construction the max-density condition holds. Finally, we prove that serving all interests $I \in \mathcal{I}^*$ with data rate at least $b(\|I\|)$ is not feasible.

Note that $b \notin o(1/d^2)$ means that there is a constant $c > 0$ such that

$$\forall d_0 > 0 : \exists d > d_0 : b(d) \geq \frac{c}{d^2}. \tag{2}$$

In the following, let c be such a constant.

Definition 5. *Let $k_0 := \max\{r_0, \Delta\}$.*

For all $i \in \mathbb{N}_{>0}$ let $k_i \in \mathbb{R}$ be a value for which

$$k_i > 8k_{i-1} + \Delta \qquad \text{and} \qquad b(k_i) \geq \frac{c}{k_i^2}.$$

Such k_i exists for all k_{i-1} because $b(d) \notin o(1/d^2)$ (to see this, set $d_0 = 8k_{i-1} + \Delta$ and $d = k_i$ in (2)).

Based on the sequence $(k_i)_{i \in \mathbb{N}}$ we can now construct the specific sets of measurement points M^* and interests $\mathcal{I}^*$ for our counterexample as follows.

Definition 6. *Let M^* be a set of measurement points defined as follows.*

We first construct a sequence of primary circles. *The primary circles are all centered at the origin $(0, 0)$. The i-th primary circle ($i \geq 0$), denoted by $C_{i,0}$, has radius k_i.*

For all $i \in \mathbb{N}_{>0}$ we construct further circles between the primary circles, in a way such that the radii of any pair of circles differ by at least Δ. We call these additional circles secondary circles. *The secondary circles, too, are centered at the origin. Between $C_{i-1,0}$ and $C_{i,0}$, there are $w_i - 1$ secondary circles, denoted by $C_{i,1}, \ldots, C_{i,w_i-1}$, where*

$$w_i = \left\lfloor \frac{k_i - k_{i-1}}{\Delta} \right\rfloor.$$

Let the radius of $C_{i,j}$ be $k_i - j\Delta$.

Figure 2 depicts the idea of primary and secondary circles (solid and dashed lines, respectively).

According to Lemma 2,

$$\left\lfloor \frac{4(k_i - j\Delta)}{\Delta} \right\rfloor$$

points can be positioned on $C_{i,j}$, such that they all have a pair-wise distance of at least Δ. M^* consists of such points on all circles $C_{i,j}$ where $i > 0$. The exact placement of the points on the circle is not relevant in the following, as long as the pair-wise distance is at least Δ; for instance, one point on $C_{i,j}$ could be positioned at coordinates $(0, k_i - j\Delta)$, with the rest of the points uniformly spaced over the circle perimeter.

Let the i-th zone*, Z_i, be the subset of M^* that contains all measurement points residing on the circles $C_{i,0}, \ldots, C_{i,w_i-1}$ (i. e., it comprises the primary circle $C_{i,0}$ and all secondary circles between $C_{i-1,0}$ and $C_{i,0}$). Observe that the zones are all of finite size, are all pairwise disjoint, and together fully cover M^*. Let z_i denote the number of measurement points in zone Z_i.*

Finally, let

$$\mathcal{I}^* = \{((0, 0), m) \mid m \in M^*\}.$$

We must now verify that a max-density condition with parameters δ and r_0 holds for our construction.

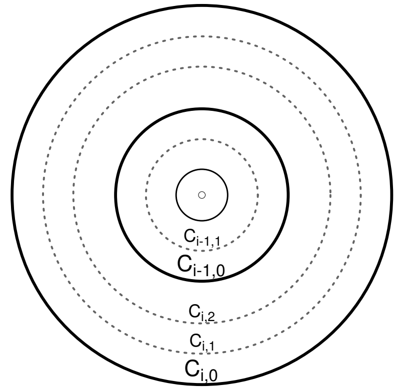

Figure 2: Primary and secondary circles in the construction of M^*.

Theorem 1. *A max-density condition with parameters δ, r_0 holds for M^*.*

Proof. We show that for any two measurement points $m_1, m_2 \in M^*, m_1 \neq m_2$ the distance $\|m_1 - m_2\|$ is at least Δ. First, note that according to Definition 6 the difference between the radii of any two circles is at least Δ. Therefore, if m_1 and m_2 reside on different circles, their distance must be at least Δ. If m_1 and m_2 are located on the same circle, however, their distance is also at least Δ by construction.

Consequently, the distance between any pair of measurement points is at least Δ. Thus, by Lemma 1, the max-density condition holds. $\square$

We now turn towards the number of measurement points within the individual zones and make the following observation regarding the number of measurement points in a zone.

LEMMA 3. *Let $i \in \mathbb{N}_{>0}$. For the number of measurement points z_i in zone Z_i it holds that*

$$z_i > \frac{k_i^2}{2\Delta^2}.$$

The lemma is proven in Appendix C.

3.4 A criterion for scalability

Finally, we are ready to prove our main theorem, which characterizes the minimum "amount" of aggregation necessary for a scalable dissemination scheme.

THEOREM 2. *Let b be an arbitrary bandwidth profile for which $b(d) \notin o(1/d^2)$. Then, for measurement points M^* and interests $\mathcal{I}^*$ as defined above, not all interests $I \in \mathcal{I}^*$ can be served with bandwidth at least $b(\|I\|)$.*

PROOF. Let M^* and $\mathcal{I}^*$ be defined like above. Consider a circle C^* centered around 0 with radius $r^* := k_0/2$. Note that, by construction, all circles in the definition of M^* have a radius larger than r^*; thus, all points in M^* are located outside C^*. Since for each $m \in M^*$ there is an interest $I_m = ((0,0), m) \in \mathcal{I}^*$, information about m must be transported into C^* at least with data rate $b(\|I_m\|)$. Note that $\|I_m\|$ is equal to the radius of the measurement point circle (according to the definition of M^*) on which m is located.

Now observe that for each measurement point m in zone Z_i, the bandwidth that must be spent for the point when delivering information for interest I_m is bounded below as follows

$$b(\|I_m\|) \geq b(k_i), \tag{3}$$

because b is monotonically decreasing per definition and $\|I_m\| \leq k_i$.

Therefore, the total rate B at which data must be transported into C^* to serve all interests according to b is—by summation over the zones—bounded below by

$$B \geq \sum_{i=1}^{\infty} b(k_i) z_i \geq \sum_{i=1}^{\infty} \frac{c}{k_i^2} z_i. \tag{4}$$

According to Lemma 3, we thus have

$$B > \sum_{i=1}^{\infty} \frac{c}{k_i^2} \left(\frac{k_i^2}{2\Delta^2} \right) = \frac{c}{2\Delta^2} \sum_{i=1}^{\infty} 1. \tag{5}$$

This sum obviously does not converge, therefore it would require infinite bandwidth into the circle C^* to serve all interests. Since the bandwidth available for transmissions into C^* is finite (bounded above by ξ_{r^*} according to Assumption 1), the assertion holds. □

Note that while from the above construction and proof it may seem that an infinite number of interests and measurement points is necessary in order to break a scheme with a bandwidth profile that is not in $o(1/d^2)$, this is not exactly true: one can always find a *finite* number of zones (and thus a finite number of measurement points and interests) for which the sum in (4) already exceeds the available bandwidth.

4. THE BOUND IS TIGHT

Let M be an arbitrary set of measurement points for which a max-density condition with arbitrary, but fixed and given parameters δ, r_0 holds. Let $\epsilon > 0$. In this section, we will argue that for any arbitrary interest set $\mathcal{I}$, data dissemination and aggregation with a bandwidth profile in $\Omega(1/d^{2+\epsilon})$ is possible. This complements the negative result from the previous section—which implies that protocols with bandwidth profiles in $\Omega(1/d^2)$ do generally not scale—with the positive counterpart that for any $\epsilon > 0$ scalable protocols with bandwidth profiles in $\Omega(1/d^{2+\epsilon})$ are feasible. We do so by showing how the available network bandwidth can be assigned to measurement points in proactive dissemination schemes, in order to achieve the desired scaling.

This bandwidth assignment is not intended to be immediately used in practice. In fact, while it might serve as a guideline for application-oriented designs in the future, due to its rather rigid structure (based on a fixed-size grid structure) it will certainly not be the best choice in typical real-world systems. Topological characteristics of the road network and other environmental factors will make higher flexibility a desirable factor. The scheme to be introduced here has the important advantage of being analytically tractable, and it can be constructed for any arbitrary $\epsilon > 0$. Both is crucial for our analytical existence proof, but will not be of equally central importance in practical protocols.

4.1 Refining the network model

Before we can tackle the proof, we need to refine our model. So far, it does not make any statement about communication that is possible in the network: Assumption 1 disallows communication with arbitrarily high rates, and thus only states what is *not* possible. But, so far, even the case where no communication is possible at all in the network is perfectly valid. Thus, before we can possibly derive a feasibility result, we need to pin down the "minimum possible communication" in the network. Following the lines of the preceding part of this paper, we again strive for a formulation that captures the essence of a broad range of networks, and therefore leads to strong and generic results. The following assumption is inspired by grid- or quadtree-based information dissemination and aggregation schemes as discussed, for example, in [3, 5, 13, 22].

ASSUMPTION 2. *We assume that there is a subdivision of $\mathbb{R}^2$ by a regular grid into square cells with side length $L > 0$, such that—once again over a sufficiently long time span—within each individual cell it is possible to distribute information to all nodes located within that cell at least with a positive data rate τ.*

The information that is distributed may stem from measurement points inside or outside the cell, in arbitrary combinations, as long as the total data rate at which information about all sources is distributed does not exceed τ. Information about a measurement point outside the cell can be distributed within a cell if and only if it is already known by the nodes in at least one neighboring cell (i. e., it is distributed there with at least the same rate).

The assumption does intentionally not specify how the dissemination is actually performed. Periodic flooding within the cells as in [22] would be one option; periodic beaconing like in [3, 13] is an alternative that is probably more robust to temporary network partitioning as it is common in VANETs. In the remainder of this section we will center our discussion around the question how the total bandwidth τ is shared between information sources.

4.2 Constructing a bandwidth assignment

In the following we will assume that $L > r_0$. We may do so without loss of generality: if this condition does not hold, we replace L repeatedly by $2L$ and the grid by a coarser one where four adjacent cells are merged into one, until $L > r_0$ holds. This does not affect Assumption 2: if Assumption 2 holds for L, τ, then it also holds for $2L$, τ, because this is equivalent to distributing the same information in the four merged cells.

DEFINITION 7. *We call the grid from Assumption 2 the* level-0 grid *and its cells* level-0 cells*. Based on the level-0 grid we construct a hierarchy of regular grids with cell side lengths $L, 2L, 4L$, etc.. Each cell of the level-i grid (termed a* level-i cell*) consists of 2×2 level-$(i-1)$ cells and has side length $2^i L$. Each cell of any lower hierarchy level lies within exactly one cell on level i.*

We now derive an upper bound on the number of measurement points within each cell from the max-density condition.

LEMMA 4. *In any level-i cell there are less than $2^{2i-1}\delta L^2$ measurement points.*

PROOF. Any level-0 cell can be fully enclosed by a circle with radius $r = L/\sqrt{2}$. Due to the max-density condition this circle contains less than $\delta r = \delta L^2/2$ measurement points. Any measurement point in the cell is also in the circle, therefore the assertion holds for level-0 cells.

A level-i cell consists of 2^{2i} level-0 cells, each containing at most $\delta L^2/2$ measurement points. Thus, the level-i cell contains no more than $2^{2i-1}\delta L^2$ measurement points and the assertion holds. □

First, we now look at the dissemination of information within the same level-0 cell. The bandwidth distribution used in our proof provides data about each measurement point to each car in the same level-0 cell at rate

$$R_{\text{local}} = \frac{\tau}{\delta L^2}.$$

Because there are at most $\delta L^2/2$ measurement points in the cell according to Lemma 4, this *local information dissemination* accounts for a data rate of at most $\tau/2$.

On each hierarchy level, aggregated (that is, bandwidth-reduced) information about the measurement points in any level-i cell C is also made available to the cars in all 9 level-i cells in a 3×3 cells block around C. According to Assumption 2 this is possible, because the information is distributed with at least the same rate in an adjacent cell. For this *aggregated information dissemination*, the data rate is reduced such that the rate at which information about a measurement point in cell C is made available on level i is

$$R_i = \frac{\tau \cdot (2^\epsilon - 1)}{18 \cdot 2^{\epsilon \cdot (i+1)} \cdot 2^{2i-1}\delta L^2}. \qquad (6)$$

Here, we are not concerned with the question exactly *how* the bandwidth reduction is performed. The cars could, for example, use techniques like those described in [3,13] to reduce the bandwidth spent for distributing information about the measurement points in larger and larger cells.

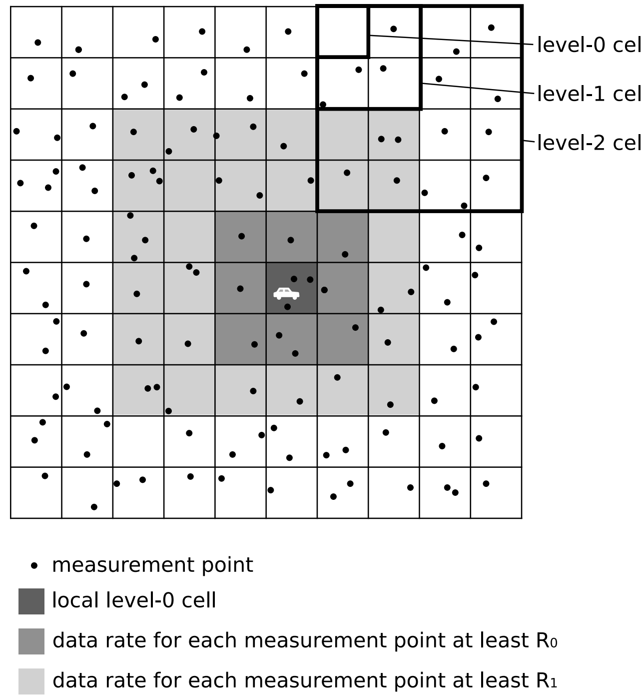

Figure 3: The hierarchical grids and the bandwidths assigned to measurement points.

The key point is that if a car is within the 3×3 cells block on level i around a measurement point, then data about this measurement point is provided to the car at least at rate R_i. The hierarchy of grids and the bandwidth assignments are visualized in Figure 3. Exemplary level-0, level-1, and level-2 cells are highlighted to show their relative sizes and the alignment of the higher-level grids. The figure also shows the position of one car in the VANET and its local level-0 cell. This car will receive information about each measurement point in the 3×3 level-0 cells block around its own local level-0 cell with a data rate of R_0. Information about each measurement point in the 3×3 level-1 cells block around the car is provided at least with rate R_1, and so on.[2]

4.3 This bandwidth assignment is scalable

It remains to show that (a) the bandwidth available according to Assumption 2 suffices, i. e., that the total rate at which information is disseminated about all measurement points does not exceed τ in any level-0 cell, and (b) that this scheme has a bandwidth profile in $\Omega(1/d^{2+\epsilon})$. We prove these two points in the following two theorems.

THEOREM 3. *The described bandwidth assignment scheme does not exceed the bandwidth limit τ in each level-0 cell.*

PROOF. Consider an arbitrary level-0 cell. On each hierarchy level i, it lies within the 3×3 block around 9 different level-i cells. For each measurement point in each of the 9 cells, information is provided to the cars within the level-0 cell at rate R_i. Recall from Lemma 4 that a level-i cell contains at most $2^{2i-1}\delta L^2$ measurement points. For hierarchy

[2]Note that the asymmetry of the 3×3 cells blocks in the figure is correct; it is a result of the relative alignments of the grids on different hierarchy levels.

level i, information is therefore distributed with a total rate of at most $9R_i 2^{2i-1}\delta L^2$. Consequently, summarizing over all hierarchy levels, the total rate used in the level-0 cell under consideration for disseminating aggregated information is no larger than

$$\sum_{i=0}^{\infty} 9R_i 2^{2i-1}\delta L^2 = \sum_{i=0}^{\infty} 9\frac{\tau\cdot(2^\epsilon-1)2^{2i-1}\delta L^2}{18\cdot 2^{\epsilon\cdot(i+1)}\cdot 2^{2i-1}\delta L^2}$$
$$= \frac{\tau(2^\epsilon-1)}{2\cdot 2^\epsilon}\sum_{i=0}^{\infty}\frac{1}{2^{\epsilon i}} = \frac{\tau(2^\epsilon-1)}{2\cdot 2^\epsilon}\cdot\frac{1}{1-2^{-\epsilon}} = \frac{\tau}{2}.$$

Since the total data rate for local information dissemination does, as noted above, also not exceed $\tau/2$, the total rate at which information must be propagated within any level-0 cell is at most τ. Thus, the described scheme does not exceed the bandwidth that is available according to Assumption 2. □

THEOREM 4. *The described bandwidth assignment scheme has a bandwidth profile in* $\Omega(1/d^{2+\epsilon})$.

PROOF. Let $I = (x, m) \in \mathcal{I}$ be an arbitrary interest with distance $d := \|I\|$. Let us first consider the case where $d > L$ and let $i := \lceil \log_2 \frac{d}{L} \rceil$. Note that then x lies within the level-i 3×3 cells block around m, because the distance between x and m is at most the side length of a level-i cell:

$$i = \left\lceil \log_2 \frac{d}{L} \right\rceil \quad\Rightarrow\quad i \geq \log_2 \frac{d}{L} \quad\Rightarrow\quad d \leq 2^i L.$$

Consequently, the interested car at position x receives information about m at a rate of at least R_i. So, if $d > L$, for the rate $B(d)$ at which information about m is provided it holds that

$$B(d) \geq R_{\lceil \log_2 \frac{d}{L}\rceil} = \frac{\tau(2^\epsilon-1)}{18\cdot 2^{\epsilon(\lceil\log_2\frac{d}{L}\rceil+1)}\cdot 2^{2\lceil\log_2\frac{d}{L}\rceil-1}\delta L^2}.$$

Because $\lceil \log_2 \frac{d}{L}\rceil < \log_2 \frac{d}{L} + 1$ it follows that

$$B(d) > \frac{\tau(2^\epsilon-1)}{18\cdot 2^{\epsilon(\log_2\frac{d}{L}+2)}\cdot 2^{2(\log_2\frac{d}{L}+1)-1}\delta L^2}$$
$$= \frac{\tau(2^\epsilon-1)}{18\cdot 2^{2\epsilon}\cdot 2^{\epsilon\log_2\frac{d}{L}}\cdot 2\cdot 2^{2\log_2\frac{d}{L}}\cdot\delta L^2}$$
$$= \frac{\tau(2^\epsilon-1)L^\epsilon}{36\cdot 4^\epsilon\cdot\delta}\cdot\frac{1}{d^{2+\epsilon}}.$$

In the case where $d \leq L$, m and x are either in the same level-0 cell, or x is in the 3×3 block of level-0 cells around m. The rate of information provisioning about m is therefore bounded below by the minimum of R_{local} and R_0. As can easily be verified, it holds that $R_0 < R_{\text{local}}$. In summary we thus have

$$B(d) \geq \begin{cases} R_0 & \text{if } d \leq L \\ \frac{\tau(2^\epsilon-1)L^\epsilon}{36\cdot 4^\epsilon\cdot\delta}\cdot\frac{1}{d^{2+\epsilon}} & \text{if } d > L. \end{cases} \tag{7}$$

Therefore, the expression on the right hand side of (7) is a bandwidth profile of the scheme constructed above. It is easy to see that this bandwidth profile is in $\Omega(1/d^{2+\epsilon})$, since only the case where $d > L$ is relevant for the asymptotic behavior and $\tau, L, \delta, \epsilon$ are all constant. □

Consequently, we can conclude that it is possible to come arbitrarily close to the limit derived in the previous section.

5. DISCUSSION

5.1 Scope and applicability

The results presented in the preceding two sections shed light on an aspect that has so far been neglected in VANET data aggregation: to which extent aggregation is needed. Of course, in practice, not only the asymptotic scaling, as it is considered here, but also the involved constants will play a significant role; moreover, one must stay aware that VANETs are huge networks, but do not have infinite extension. Therefore, in practice, careful consideration remains inevitable, and application-specific aspects must be taken into account. VANET applications will be long-lived deployments, where the introduction of revised or even re-designed protocols is difficult or impossible. At design time, the detailed characteristics of all (road) networks in which an application will be used will typically not be known; the same holds true for the future usage pattern of an application. The notion of a bandwidth profile and the results on the asymptotic scaling of aggregation schemes therefore constitute valuable guidelines when mechanisms for VANET dissemination applications are selected, adapted, or designed: such mechanisms should at least *allow* to avoid bandwidth profiles which are not in $o(1/d^2)$—otherwise they are bound to fail in the general case. Some, but not all of the existing mechanisms already provide the necessary flexibility. In any case, before aggregation mechanisms are considered as a basis for large-scale applications, they should be subjected to a rigorous analysis with respect to their bandwidth requirements and their scaling behavior.

We introduced and motivated our results with dissemination applications in VANETs. However, due to the generality of the assumptions on which our proofs are based, our results are not limited to this application field; they likewise hold for many other environments. Nevertheless, we believe that they are much more relevant in the VANET context than in other settings: in VANET data dissemination applications, the nodes continuously measure different parameters in a larger geographic area, aggregate these measurements within the network, and communicate the aggregated measurements to many interested parties in different areas. This is very different to, for instance, sensor networks, where the data are reported to typically very few sinks. It is this difference that makes data aggregation in VANETs so challenging.

5.2 An example

Out of the many proposed aggregation schemes, let us pick SOTIS [24] and the aggregation mechanism that we proposed in [13] in order to show what our results mean in practice. Both approaches use periodic beacons for proactive data dissemination. If a vehicle receives such a beacon, it merges the contained information into a local knowledge base, and sends out subsets of the information from its knowledge base in its own beacons.

Where SOTIS uses fixed-size road segments to summarize travel speeds, we show in [13] how (more generic) data can be increasingly aggregated over larger and larger areas with increasing distance. The focus of that work is on the data representation aspect of aggregation: how measured parameters should be encoded in order to allow for data from many sources to be collected and aggregated in a fully distributed setting. The example application described in [13] measures and distributes information on available parking slots. The reference implementation used there is based on a grid structure quite similar to the one described in the proof in Section 4: for each level of aggregation four grid elements are aggregated to form one grid element of the next higher level.

Given its own position, each car then has a set of data elements that it can choose to report on. This could be single road segments in SOTIS, or grid cells on any hierarchy level for the hierarchical aggregation scheme. The car must decide on the subset of elements about which information is included in the beacons, and also *how often* it is included—in every beacon, in every other beacon, etc.. The choice can be captured by defining a (position-dependent) frequency f_E for each element E of the set, where a car includes information about E in its beacons with frequency f_E. These frequencies implicitly also determine the beaconing interval and the size of the beacons, and therefore the network bandwidth used.

For such a model, our previous results are directly applicable. If we assume, for simplicity of discussion, a constant density of measurement points and a constant aggregate size (both seems reasonable for both discussed schemes), we can obtain results on how we must choose the frequencies. For the hierarchical aggregation scheme, we can build upon the proof from the previous section and start from the rates per measurement point in (6). Just as in the proof, we may select all cells in the 3×3 cells block around the sender's current position for transmission, on each hierarchy level i. If all parameters except i are fixed (including fixed, but arbitrary $\epsilon > 0$), then it is easy to see that the rate per measurement point for level i is in $\Theta(1/(2^{\epsilon i} \cdot 2^{2i}))$.

With the assumption of constant measurement point density, we have $\Theta(2^{2i})$ measurement points in a level-i cell. We can therefore afford a total rate in $\Theta(1/2^{\epsilon i})$ for the entirety of measurement points in a level-i cell—i. e., for the aggregate describing that cell. While the exact constants will, of course, depend on the specific setting, this gives us a pretty clear idea about how often we may transmit information on a region: we should decrease the frequency with increasing aggregation level, but we may do so very slowly (because ϵ can be arbitrarily small). Then, the arguments above are immediately transferrable, which proves scalability and a reasonable bandwidth profile.

For SOTIS, where no hierarchical aggregation is used, the picture is a bit different. We still assume constant measurement point density (in this case, a constant number of road segments per map area). From the results in Section 3, we learn that this can only be scalable if we reduce the rate like $o(1/d^2)$ over distance. With constant aggregate size, we would indeed need a quite rapid reduction of the frequency to achieve the necessary rate reduction—faster than $1/d^2$. When the transmission frequency is reduced so quickly for more distant regions, however, it might be questionable whether these rare transmissions of information about more distant segments can be practically useful at all: cars might barely be able to obtain useful information about more distant regions within reasonable time spans.

5.3 The impact of mobility

We believe that one aspect deserves further discussion with respect to the applicability of our model to practical VANETs: the well-known result by Grossglauser and Tse [6], who showed that node mobility can increase the capacity of mobile wireless networks if nodes can "carry" data with them, passing the information on to other nodes at later times. The delay-tolerant communication approaches inspired by these results are vital for bridging connectivity gaps in the network, and are thus a central building block of many VANET protocols. However, as has been pointed out before for example in [16,18], the assumptions that lead to this insight—unlimited buffers, (arbitrarily) high delay tolerance and independent node movement—do not apply in practical VANETs. If finite buffers are assumed, a basic observation confirms our conjecture that, in practice, mobility will not fundamentally affect the validity of Assumption 1: just like the total bandwidth of the communication links crossing the boundary of a given finite-size circle, the number of vehicles per time unit crossing the circle boundary (each carrying a limited amount of information) is also inherently limited. Mobility may thus impact the value of ξ_r in Assumption 1, but not the assumption's fundamental validity.

6. RELATED WORK

So far, aggregation in VANETs has been considered mostly from a practical perspective. Early works in the VANET area that considered aggregation as we understand it here—reducing the amount of distributed information at larger distances—were the distributed traffic information systems SOTIS by Wischhof et al. [24] and TrafficView by Nadeem et al. [15]. Like in the later proposed approach CASCADE by Ibrahim and Weigle [8], the main focus of these systems is on disseminating information along one single road; consequently, the proposed aggregation mechanisms aim at combining information from measurement points or vehicles on the same road. Raya et al. [20] consider the secure aggregation of messages in VANETs, and Yu et al. [26] discuss how reports generated by multiple cars—again generated on one single road (segment)—can efficiently be collected in order to generate aggregate reports.

The probably first protocol that performs hierarchical in-network data aggregation over a whole city road network is the parking guidance system proposed by Caliskan et al. [3]. Dornbush and Joshi [4] introduce StreetSmart, a scheme that applies data mining techniques in the cars in order to extract the most relevant features that characterize the current traffic situation in a network of roads. We discuss a data representation for area-based hierarchical aggregation using modified FM-sketches in [13] and show how hierarchical aggregation can be applied to travel times in road networks [14]. All these references introduce specific mechanisms and data structures that can be used to actually perform aggregation, but none of them reflects on how these schemes should be applied and parameterized in order to yield a scalable overall system. Our results presented here are thus complementary: while generic enough to be applicable to all these aggregation schemes, they are the first results that provide insight on the extent to which distance-based in-network aggregation is necessary and useful.

Data aggregation has also been considered in the context of wireless sensor networks [1, 5, 12, 19]. There, the main focus lies on reducing the energy consumption when transmitting information to one or few sinks. In VANETs, energy is not an issue, and information must continuously be provided to most or all network participants. Thus, the questions that need to be solved are very different.

The feasibility and infeasibility considerations in this paper are in some aspects related to existing results on the capacity of wireless multihop networks. A large body of theoretical work followed the milestone paper by Gupta and Kumar [7], who introduced a framework that has subsequently often been used. Their results on unicast communication have later been generalized to a broader class of communication patterns; for instance, Keshavarz-Haddad et al. [9] assess the capacity for broadcast communication, and Shakkottai et al. [21] consider multicast. Wang et al. [23] present generic results for a broader family of communication paradigms termed (n, m, k)-casting. In [10], results for unicast, multicast and broadcast are derived based on a model that differs in some fundamental aspects from the one introduced by Gupta and Kumar. An important aspect in the context of our work is the distance between sender and receiver, which is taken into account in the distance-weighed throughput metric used in [25]. For the specific case of VANETs, Pishro-Nik et al. [18] assessed capacity scaling laws for unicast communication and their dependency on the characteristics of the road network; within the same framework, distance-limited unicast and broadcast communication is analyzed in [16]. None of these works, however, is applicable to the questions posed here, because none of them considers any form of in-network aggregation.

As already mentioned, Grossglauser and Tse [6] and some subsequent works showed that—theoretically—the capacity of wireless networks increases if node mobility is taken into account. However, as discussed in the preceding section, mobility does not fundamentally affect our results. From a more information-theoretic perspective and with models that are less technology- and protocol-centric than the ones in most works cited above, it has been shown that higher capacities can be achieved, if cooperative transmissions and techniques like distributed MIMO are considered; see, e. g., [17, 25]. However, under realistic assumptions about what can be realized in real-world VANETs, this, too, does not fundamentally change the validity of our assumptions about possible and impossible communication, and therefore also does not affect our results.

7. CONCLUSION

In this paper we have analyzed information dissemination in the context of vehicular networks. Due to capacity limitations, data aggregation techniques that reduce the amount of information with increasing distance are vitally necessary for dissemination-based applications in VANETs. This paper contains two central contributions in this context: first, we introduced the concept of a bandwidth profile, and showed that aggregation schemes cannot be considered generally scalable unless their bandwidth profiles are in $o(1/d^2)$, where d is the geographical distance between a source of information and an interested vehicle. Second, by constructing a respective bandwidth distribution for proactive information dissemination, we proved that the this bound is tight: for arbitrary $\epsilon > 0$, a bandwidth profile in $\Omega(1/d^{2+\epsilon})$ can be achieved.

Acknowledgments

The authors thank the German Research Foundation (DFG) for the support.

8. REFERENCES

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: A survey. *Elsevier Computer Networks*, 38(4):393–422, Mar. 2002.

[2] S. Basagni, I. Chlamtac, V. R. Syrotiuk, and B. A. Woodward. A distance routing effect algorithm for mobility (DREAM). In *MobiCom '98: Proceedings of the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking*, pages 76–84, Oct. 1998.

[3] M. Caliskan, D. Graupner, and M. Mauve. Decentralized discovery of free parking places. In *VANET '06: Proceedings of the 3rd ACM International Workshop on Vehicular Ad Hoc Networks*, pages 30–39, Sept. 2006.

[4] S. Dornbush and A. Joshi. StreetSmart Traffic: Discovering and disseminating automobile congestion using vanet's. In *VTC '07-Spring: Proceedings of the 65th IEEE Vehicular Technology Conference*, pages 11–15, Apr. 2007.

[5] M. Enachescu, A. Goel, R. Govindan, and R. Motwani. Scale-free aggregation in sensor networks. *Elsevier Theoretical Computer Science*, 344(1):15–29, Nov. 2005.

[6] M. Grossglauser and D. N. C. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking*, 10(4):477–486, 2002.

[7] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Transactions on Information Theory*, IT-46(2):388–404, Mar. 2000.

[8] K. Ibrahim and M. C. Weigle. Optimizing CASCADE data aggregation for VANETs. In *MoVeNet '08: Proceedings of the 2nd International Workshop on Mobile Vehicular Networks*, Sept. 2008.

[9] A. Keshavarz-Haddad, V. Ribeiro, and R. Riedi. Broadcast capacity in multihop wireless networks. In *MobiCom '06: Proceedings of the 12th Annual ACM International Conference on Mobile Computing and Networking*, pages 239–250, Sept. 2006.

[10] A. Keshavarz-Haddad and R. Riedi. Bounds for the capacity of wireless multihop networks imposed by topology and demand. In *MobiHoc '07: Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 256–265, Sept. 2007.

[11] T. Kosch, C. Schwingenschlögl, and L. Ai. Information dissemination in multihop inter-vehicle networks. In *ITSC '02: Proceedings of the 5th International IEEE Conference on Intelligent Transportation Systems*, pages 685–690, Sept. 2002.

[12] B. Krishnamachari, D. Estrin, and S. Wicker. The impact of data aggregation in wireless sensor networks. In *ICDCSW '02: Proceedings of the 22nd International Conference on Distributed Computing Systems Workshops*, pages 575–578, July 2002.

[13] C. Lochert, B. Scheuermann, and M. Mauve. Probabilistic aggregation for data dissemination in VANETs. In *VANET '07: Proceedings of the 4th ACM International Workshop on Vehicular Ad Hoc Networks*, pages 1–8, Sept. 2007.

[14] C. Lochert, B. Scheuermann, C. Wewetzer, A. Luebke, and M. Mauve. Data aggregation and roadside unit placement for a VANET traffic information system. In *VANET '08: Proceedings of the 5th ACM International Workshop on VehiculAr Inter-NETworking*, pages 58–65, Sept. 2008.

[15] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode. TrafficView: traffic data dissemination using car-to-car communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 8(3):6–19, July 2004.

[16] M. Nekoui, A. Eslami, and H. Pishro-Nik. Scaling laws for distance limited communications in vehicular ad hoc networks. In *ICC '08: Proceedings of the IEEE International Conference on Communications*, pages 2253–2257, May 2008.

[17] A. Özgür, O. Lévêque, and D. N. C. Tse. Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks. *IEEE Transactions on Information Theory*, 53(10):3549–3572, 2008.

[18] H. Pishro-Nik, A. Ganz, and D. Ni. The capacity of vehicular ad hoc networks. In *Allerton '07: Proceedings of the 45rd Annual Allerton Conference*, pages 1156–1163, Sept. 2007.

[19] R. Rajagopalan and P. K. Varshney. Data aggregation techniques in sensor networks: A survey. *IEEE Communications Surveys & Tutorials*, 8(4):48–63, 2006.

[20] M. Raya, A. Aziz, and J.-P. Hubaux. Efficient secure aggregation in VANETs. In *VANET '06: Proceedings of the 3rd ACM International Workshop on Vehicular Ad Hoc Networks*, pages 67–75, Sept. 2006.

[21] S. Shakkottai, X. Liu, and R. Srikant. The multicast capacity of large multihop wireless networks. In *MobiHoc '07: Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 247–255, Sept. 2007.

[22] M. Transier, H. Füßler, J. Widmer, M. Mauve, and W. Effelsberg. A hierarchical approach to position-based multicast for mobile ad hoc networks. *Springer Wireless Networks*, 13(4):447–460, Aug. 2007.

[23] Z. Wang, H. R. Sadjadpour, and J. J. Garcia-Luna-Aceves. A unifying perspective on the capacity of wireless ad hoc networks. In *INFOCOM '08: Proceedings of the 27th Annual Joint Conference of the IEEE Computer and Communications Societies*, pages 753–762, Apr. 2008.

[24] L. Wischhof, A. Ebner, H. Rohling, M. Lott, and R. Halfmann. SOTIS – a self-organizing traffic information system. In *VTC '03-Spring: Proceedings of the 57th IEEE Vehicular Technology Conference*, pages 2442–2446, Apr. 2003.

[25] L.-L. Xie and P. R. Kumar. A network information theory for wireless communication: Scaling laws and optimal operation. *IEEE Transactions on Information Theory*, 50(5):748–767, May 2004.

[26] B. Yu, J. Gong, and C.-Z. Xu. Catch-up: A data aggregation scheme for VANETs. In *VANET '08: Proceedings of the 5th ACM International Workshop on VehiculAr Inter-NETworking*, pages 49–57, Sept. 2008.

APPENDIX

A. PROOF OF LEMMA 1

We first note that Δ is well-defined and positive because from $\delta r_0^2 > 1$ if follows that $\sqrt{\delta} - r_0^{-1} > 0$. We must show that for any circle C with radius $r \geq r_0$ less than δr^2 points can be fit into C, if the pair-wise distances between the points are all at least Δ. The latter is equivalent to stating that if we draw circles with radius $\frac{\Delta}{2}$ around each of the measurement points within C, then none of these small circles may intersect with each other.

Observe that the centers of all small circles must lie within C. Therefore, all the small circles are fully enclosed within a circle C^+ with the same center as C and radius $r + \frac{\Delta}{2}$. Thus, the total area of the small circles (each covering an area of $\pi(\Delta/2)^2$) is less than the area of C^+, which is $\pi(r + \Delta/2)^2$. (It is strictly less because C^+ can impossibly be fully covered by small circles.) Therefore, for the total number n of small circles (i. e., the total number of measurement points within C) it holds that

$$n < \frac{\pi\,(r + (\Delta/2))^2}{\pi(\Delta/2)^2} = \frac{r^2}{(\Delta/2)^2} + \frac{2r}{\Delta/2} + 1$$
$$= r^2\left(\sqrt{\delta} - \frac{1}{r_0}\right)^2 + 2r\left(\sqrt{\delta} - \frac{1}{r_0}\right) + 1.$$

Because $r \geq r_0$, n is thus bounded above by

$$r^2\left(\sqrt{\delta} - \frac{1}{r}\right)^2 + 2r\left(\sqrt{\delta} - \frac{1}{r}\right) + 1 = r^2\delta.$$

Consequently, the max-density condition holds. $\square$

B. PROOF OF LEMMA 2

We show that when $\left\lfloor \frac{4r}{\Delta} \right\rfloor$ points are evenly distributed on the perimeter of a circle with radius $r \geq \frac{\Delta}{2}$, the distance between two neighboring points is at least Δ. The angle α between two neighboring such points on the circle is

$$\alpha = \frac{2\pi}{\left\lfloor \frac{4r}{\Delta} \right\rfloor} \geq \frac{2\pi}{\frac{4r}{\Delta}} = \frac{\pi}{2}\frac{\Delta}{r}.$$

The distance μ between the points is the base of an isosceles triangle with arm length r and angle α. Thus,

$$\mu = 2r\sin\left(\frac{\alpha}{2}\right).$$

Note that $\forall x \in [0, \pi/2]$ it holds that $\sin(x) \geq \frac{2x}{\pi}$, and also note that $\alpha < \pi$ (because $r \geq \frac{\Delta}{2}$ and thus $\left\lfloor \frac{4r}{\Delta} \right\rfloor \geq 2$). Therefore,

$$\mu \geq 2r\frac{2\alpha}{\pi} \geq 2r\frac{2}{\pi}\frac{\pi\Delta}{4r} = \Delta.$$

From the fact that the distance between each pair of neighboring points is at least Δ, it easily follows that all pair-wise distances are at least Δ. $\square$

C. PROOF OF LEMMA 3

First, we show a simple lemma which will be of great help.

LEMMA 5. *For all $k \in \mathbb{N}, k \geq 1$ and all $x \geq k$ it holds that*

$$\lfloor x \rfloor > \frac{k}{k+1}x.$$

PROOF. For $x \in \mathbb{N}$ the assertion trivially holds. Thus, we focus on the case $x \notin \mathbb{N}$. Let $y := x - \lfloor x \rfloor$. Since $k \in \mathbb{N}$ and $x \geq k$ we also have $\lfloor x \rfloor \geq k$. Because $0 < y < 1$ we get

$$\frac{\lfloor x \rfloor}{y} > k \quad \Rightarrow \quad \lfloor x \rfloor > ky \quad \Rightarrow \quad \lfloor x \rfloor > k(x - \lfloor x \rfloor)$$
$$\Rightarrow \quad (k+1)\lfloor x \rfloor > kx \quad \Rightarrow \quad \lfloor x \rfloor > \frac{k}{k+1}x.$$

□

We now turn towards proving Lemma 3.

According to the definition of M^* above, the total number of measurement points in zone Z_i is

$$z_i = \sum_{j=0}^{w_i-1} \left\lfloor \frac{4(k_i - j\Delta)}{\Delta} \right\rfloor$$

Since $\frac{4(k_i - j\Delta)}{\Delta} > 1$ for any $j < w_i$, we can apply Lemma 5 (with $k = 1$) and get

$$\begin{aligned} z_i &> \sum_{j=0}^{w_i-1} \frac{2(k_i - j\Delta)}{\Delta} \\ &= 2\sum_{j=0}^{w_i-1} \frac{k_i}{\Delta} - 2\sum_{j=0}^{w_i-1} j \\ &= 2\left\lfloor \frac{k_i - k_{i-1}}{\Delta} \right\rfloor \frac{k_i}{\Delta} - 2\frac{(w_i-1)w_i}{2}. \end{aligned}$$

Since per definition $k_i > 8k_{i-1}$ and $\forall i \in \mathbb{N} : k_i \geq \Delta$ it holds that

$$\frac{k_i - k_{i-1}}{\Delta} > 7.$$

We may thus again apply Lemma 5 (this time with $k = 7$) and obtain

$$z_i > \frac{7(k_i^2 - k_i k_{i-1})}{4\Delta^2} - (w_i - 1)w_i.$$

Recall that the number of circles in zone Z_i is

$$w_i = \left\lfloor \frac{k_i - k_{i-1}}{\Delta} \right\rfloor$$

and therefore

$$1 \leq w_i < \frac{k_i}{\Delta}.$$

Thus,

$$(w_i - 1)w_i < \frac{k_i^2}{\Delta^2}$$

and we arrive at

$$\begin{aligned} z_i &> \frac{7(k_i^2 - k_i k_{i-1})}{4\Delta^2} - \frac{k_i^2}{\Delta^2} \\ &= \frac{3}{4}\frac{k_i^2}{\Delta^2} - \frac{7}{4}\frac{k_i k_{i-1}}{\Delta^2} \\ &= \frac{1}{2}\frac{k_i^2}{\Delta^2} + \left(\frac{1}{4}\frac{k_i^2}{\Delta^2} - \frac{7}{4}\frac{k_i k_{i-1}}{\Delta^2}\right) \\ &= \frac{1}{2}\frac{k_i^2}{\Delta^2} + \left(\frac{(k_i - 7k_{i-1})k_i}{4\Delta^2}\right). \end{aligned}$$

Since per definition $k_i > 7k_{i-1}$, the term in parentheses is positive and thus

$$z_i > \frac{k_i^2}{2\Delta^2}.$$

This is the assertion. □