



HAL
open science

Vehicular networking: from fundamental properties to network solutions

Marco Fiore

► **To cite this version:**

Marco Fiore. Vehicular networking: from fundamental properties to network solutions. Networking and Internet Architecture [cs.NI]. INSA Lyon; Université Claude Bernard Lyon 1, 2014. tel-01090951

HAL Id: tel-01090951

<https://inria.hal.science/tel-01090951>

Submitted on 4 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



A thesis submitted in fulfillment of the requirements for the degree of

Habilitation à diriger des recherches

by

Marco Fiore

Vehicular networking: from fundamental properties to network solutions

July 23, 2014

Committee

Reviewers	André-Luc Beylot	Full Professor, ENSEEIHT
	Marcelo Dias de Amorim	Research Director, CNRS
	Pietro Manzoni	Full Professor, UPV
Examiners	Carla-Fabiana Chiasserini	Associate professor, Politecnico di Torino
	Andrzej Duda	Full Professor, Grenoble INP – Ensimag
	Isabelle Guérin-Lassous	Full Professor, Université Lyon I
	Fabrice Valois	Full Professor, INSA de Lyon

A te che sei così Piccola e così grande.

Acknowledgments

If I managed to write this manuscript, most of the credit goes to all the people that, every day, indirectly support or directly participate in my work. To them all, my sincere thanks.

I owe special thanks to: Fabrice, for his tireless and precious advising throughout the past five years; Razvan, for his brilliance, his love for football and his *canapè*; Sandesh and Diala, for their hard work and for bearing with me being a perfectionist; Oscar, for his extreme efficiency and for never complaining at my bizarre requests for additional simulations; Jérôme, for his commitment to research, his friendship and his 31 year old wine; Carla and Claudio, for teaching me how to do research and for their continuous support; Francesco, José Maria, Panos, and Carlo, for being great co-authors; the members of the committee, who accepted to participate in my *Habilitation* and read this 120-page document; all the members of the UrbaNet team, for their constant attempts at redefining the concept of efficiency in teamworking; all the members of the CITI laboratory, for their smiling attitude and their passion for *Tiramisù*; all the members of the Télécom Department of INSA Lyon, who showed me how teaching can be fun for both professors and students; my whole family, for every aspect of my life that does not concern research (luckily, there are some).

Contents

1	Introduction	1
1.1	Foreword	1
1.2	Early research topics	2
1.2.1	Information management in VANETs	3
1.2.2	Mobility in vehicular networks	3
1.2.3	Vehicular handover in urban-scale mesh networks	4
1.3	Evolution of the research context	4
1.4	Selected research results	6
1.4.1	Scaling up mobility modeling and assessing its impact	6
1.4.2	From VANET traffic to data management for networked cars	7
1.4.3	The need for secure positioning	8
1.5	Document contents and structure	9
	Bibliography	9
2	Fundamental properties of vehicular networking	13
2.1	Modeling road traffic for vehicular networking	14
2.1.1	State of the art	14
2.1.2	The Cologne dataset	16
2.1.3	Relevance of road traffic realism	19
2.2	Instantaneous topology of a vehicular network	21
2.2.1	Network level	21
2.2.2	Component level	24
2.2.3	Node level	25
2.2.4	Discussion	26
2.3	Pervasive vehicular access	27
2.3.1	Access load volumes	28
2.3.2	Access load flows	30
2.3.3	Cell inter-arrival times	32
2.3.4	Cell residence times	33
2.4	Conclusions of the Chapter	35
	Bibliography	35
3	Data transfers in vehicular networks	39
3.1	The downlink direction: cooperative download	40
3.1.1	Drawing the performance bounds	41

3.1.2	Practical solutions	48
3.1.3	Discussion	54
3.2	The uplink direction: floating car data collection	55
3.2.1	Drawing the performance bounds	57
3.2.2	Practical solutions	60
3.3	Conclusions of the Chapter	63
	Bibliography	64
4	Secure positioning in vehicular environments	66
4.1	Distributed secure positioning	67
4.1.1	NPV protocol operation	67
4.1.2	NPV protocol performance evaluation	69
4.2	Centralized secure positioning	71
4.2.1	A-VIP framework operation	72
4.2.2	A-VIP framework performance evaluation	76
4.3	Conclusions of the Chapter	77
	Bibliography	78
5	Perspectives	79
5.1	Medium-term interests	80
5.1.1	Vehicular networking	81
5.1.2	Mobile data analysis	84
5.2	Long-term perspectives	88
	Bibliography	89
A	Curriculum Vitae	92

List of Figures

2.1	Cologne mobility dataset features. (a) Road network layout. (b) Road traffic snapshot at 7:00 am. (c) Simulated vehicles status over time.	16
2.2	Cologne mobility dataset generation workflow.	17
2.3	Cologne mobility dataset validation. (a) Real-world traffic at 5:00 pm. (b) Synthetic road traffic at 5:00 pm. (c) Measurement locations of road traffic. (d) Daily traffic volumes at each such location.	18
2.4	Comparative evaluation of the vehicular network performance in presence of different mobility datasets. (a) Road traffic over time. (b) Epidemic dissemination ratio over time with a penetration rate of 1.0. (c) Latency to achieve the dissemination ratio quantiles 0.05, 0.25, 0.5, 0.75 and 0.95 under different penetration rates.	19
2.5	Network-level metrics. (a) Aggregate distribution of the number of components, \mathcal{C} . (b) Aggregate distribution of the size of components, \mathcal{S} . (c) Network availability over the day.	22
2.6	Spatiotemporal view of the network-level connectivity. Each plot is a geographical representation of the network components at a specific time of the day. This figure is best viewed in colors.	23
2.7	Component-level metrics. (a) Small world property of the vehicular network versus the number of graph nodes, \mathcal{N} . (b,c) Internal structure of the largest component at two subsequent time seconds, where vertices color and size indicate the node's betweenness centrality.	24
2.8	Node-level metrics, aggregated over the whole day. (a,b)) CDF and CCDF of the vertex degree \mathcal{D} . (c) Vertex degree assortativity.	25
2.9	(a)–(m) Spatiotemporal evolution of the expected data traffic load generated by pervasive vehicular access in the Cologne region, during one day. (n) Typical daily profile of normalized RAN traffic load – concession of the Autonomous Networks Research Group at USC, http://anrg.usc.edu . (o) Geographical distribution of the population in the Cologne region. This figure is best viewed in colors.	29
2.10	Spatiotemporal evolution of macroscopic flows of pervasive vehicular access in the Cologne region over one day. This figure is best viewed in colors.	31

2.11	Distributions of vehicular user inter-arrival time at RAN cells. (a) Aggregate distributions over 30-minute intervals. (b,c) Theoretical fittings on the distribution recorded at 8:00 pm, in linear-linear and logarithmic-logarithmic scales.	32
2.12	Distributions of vehicular user residence time within RAN cells. (a) Aggregate distributions over 30-minute intervals. (b,c) Theoretical fittings on the distribution recorded at 8:00 pm, in linear-linear and logarithmic-logarithmic scales.	34
3.1	Cooperative download scenario involving DSRC-enabled vehicles. . .	40
3.2	Dynamic network topology graph (DNTG) sample. (a) Contact events among vehicles and one candidate RSU location, with shadowed areas representing halved transmission ranges, so that links exist when two shadowed areas touch or overlap, and break when such areas become disjoint. (b) Corresponding DNTG.	42
3.3	Cooperative download: average per-downloader metrics obtained from the max-flow solution, imposing a 2-hop limit to relaying. Results are shown as a function of the technology penetration rate. (a) Throughput. (b) Delay. (c) V2V downloading fraction.	44
3.4	Cooperative download at low-penetration regime. (a) Average per-downloader throughput. (b) Number of transfer hops split-up of the average per-downloader throughput, when the number of relays is unconstrained.	45
3.5	Cooperative download at low-penetration regime. (a) Transfer paradigm split-up of the average per-downloader throughput. (b) Average per-downloader throughput versus the distance between the downloader and the closest RSU.	46
3.6	Cooperative download at high-penetration regime, versus the fraction of concurrent downloaders. (a) Average per-downloader throughput. (b) Fairness.	47
3.7	Cooperative download at high-penetration regime, in diversely urbanized areas. (a) Average per-downloader throughput. (b) V2V downloading fraction.	47
3.8	Cooperative download via a probabilistic graph-based scheduling, versus the vehicular mobility prediction accuracy. (a) Fraction of demanded content that is retrieved through cooperative download. (b) Efficiency of the cooperative download.	49
3.9	Production and forward phase notation.	51
3.10	Pseudocode for relay selection at RSU B	52
3.11	Cooperative download via a contact map-based scheduling. (a) Working points of the different relay selection algorithms in the space of download rate and undelivery ratio. (b) Number of relays carrying a same content chunk for different relay selection schemes, in different road traffic scenarios.	54
3.12	FCD upload scenarios. (a) Traditional via the cellular infrastructure. (b) Offloaded through V2V transfers.	57

3.13 Offloaded FCD upload performance bounds. (a) CDF of the optimal system gain over 24 hours. (b) Dynamics of the optimal system gain over time. (c) Typical normalized traffic load at the cellular access network, over a 24-hour period. The latter data was obtained from experiments conducted by the Autonomous Networks Research Group at USC, <http://anrg.usc.edu>. 58

3.14 Spatial dynamics of offloaded FCD upload, obtained by separating the gain on the 86 *Stadtteile*, i.e., districts, of Cologne. (a) System gain at 7.30 am. (b) System gain at 12.30 am. (c) System gain at 9 pm. Figure best viewed in colors. 59

3.15 Practical schemes for offloaded FCD upload. (a) Optimal k and minimal cost of PR-C . (b) Performance of the PR-C and RB mechanisms during the morning road traffic peak. 62

4.1 Neighbor position verification (NPV) protocol operation. (a) Message exchange overview during one instance of the protocol. (b) Example of topological information stored by verifier S at the end of the message exchange and effect of a fake position announcement by M 68

4.2 Neighbor position verification (NPV) protocol performance. Probability of incorrect classification versus the ratio of vehicles that are adversarial. (a) False negatives and false positives. (b) Unverifiables. 70

4.3 Neighbor position verification (NPV) protocol performance. Probability of incorrect classification versus the number of adversaries that collude in running each attack. (a) False negatives and false positives. (b) Unverifiables. 70

4.4 A-VIP procedures by beaconer, reporter and LA. 72

4.5 A-VIP position estimation example. (a) The shaded area represents the transmission range of v_i , while the annuluses denote the set of locations from which a beacon could be received by, respectively, v_k and v_j with the quality level indicated in their report. (b) the intersection of the annuluses represents the possible positions of v_i according to the A-VIP position estimation process on reports by v_k and v_j 74

4.6 A-VIP performance evaluation, in presence of false location attacks. (a) Simulation: distribution of the trustworthiness probability γ , with 10% of adversaries announcing 10 s-old positions and a beaconing interval τ_b equal to 1 or 10 s. (b) Simulation: false positives and negatives versus fraction of attackers, with 10% of adversaries announcing 10 s-old positions and a beaconing interval τ_b of 10 s. (c) Testbed: error of the final positions determined by the LA, with 2 attackers out of 6 vehicles announcing 10 s- or 30 s-old positions and a beaconing interval τ_b of 1 s. 77

5.1 Induction loop (a). Close-by view of two measurement points (b,c). Geographical location of the two points in the conurbation of Madrid, Spain (d). 82

5.2	Distributions of the number of components \mathcal{C} (top), and of the largest component size \mathcal{S}_{max} (bottom), for each trace under different values of R	83
5.3	Workflow of the framework for the definition of categories of network usage profiles. Phase 1: construction of the snapshot graph from snapshots (portrayed here as geographical plots of the mobile traffic volume) in the training set. Phase 2: iterative aggregation of graph vertices into a dendrogram structure. Phase 3: identification of the clustering level n granting the maximum separation between the groups of snapshots. The resulting clusters are mapped to network usage profile categories.	85
5.4	Classification of two weeks of snapshots into mobile traffic profile categories, according to two different similarity measures. Each square represents a 1-hour snapshot, whose category maps to a color. Each row maps to one day. Empty squares are hours for which data was not available.	86
5.5	The complementary cumulative distribution function (CCDF) of the edge persistence (a,c) and topological overlap (b,d) for the real-world contact graph and its random equivalent. Results refer to two different four-week datasets, collected at Dartmouth (a,b) and USC (c,d) campuses.	87
5.6	Snapshots of the Dartmouth (a,b) and USC (c,d) contact graphs, considering only the social (a,c) edges or random (b,d) edges.	88

List of Tables

2.1 Major features of the vehicular mobility traces currently available for network simulation	15
--	----

Chapter 1

Introduction

1.1 Foreword

This document provides an overview of the research activities I carried out on different, yet intertwined, aspects relating to the field of vehicular networking. It covers my work over the past six years – from early 2008, when I received my PhD degree from Politecnico di Torino, Italy, until late 2013.

During this period I held different positions at multiple institutions in Europe. Between April 2008 and August 2009, I was a postdoctoral fellow at Politecnico di Torino, Italy, and I spent three months, at the end of 2008, as a visiting scholar at Universitat Politecnica de Catalunya, Spain. In September 2009 I moved to the Institut National des Sciences Appliquées (INSA) de Lyon, France, as a Maître des Conférences, the equivalent of a tenured assistant professor¹. Since my arrival at INSA Lyon, I have also become part of Inria as a member of different research teams² hosted by the CITI laboratory. Specifically, until December 2011, I was a member of the SWING team, which focused on cross-layer solutions for future wireless networks. In January 2012, I participated, jointly with Dr. Hervé Rivano and Prof. Fabrice Valois, in the creation of a brand new team, UrbaNet, which focuses on capillary³ wireless access network infrastructures for smart cities. Finally, in March 2013, I moved to Consiglio Nazionale delle Ricerche (CNR), Italy, within the Istituto di Elettronica ed Ingegneria Informatica e delle Telecomunicazioni (IEIIT), where I am currently a researcher – although I still hold an associate researcher position within the Inria UrbaNet team. Throughout all these years, I obtained most of my results within the framework of collaborations, with my former PhD thesis advisors

¹Or of an associate professor, depending on one's experience (or ego).

²Inria (formerly known as INRIA, i.e., Institut National pour la Recherche en Informatique et Automatique, before marketing experts deemed that acronyms are not sexy enough) is structured into a limited number of small teams of researchers working together on a well-defined topic, for a period of four to eight years.

³The term *capillary* provides an intuitive picture of the kind of networks that are the focus of UrbaNet: similar to capillaries in the body network of blood vessels, the wireless networks studied by the team represent the very last hop between the end user and the rest of the telecommunication network in smart cities. The analogy is opportune also when considering that capillary networks are expected to be pervasive in coverage and highly distributed from a geographical viewpoint. For more details please refer to [1].

at first, and with colleagues or advised master and PhD students later on. Their high-quality, hard work resulted in a significant body of publications, covering rather heterogeneous topics in the broad field of mobile networking.

The main common thread of such activities has remained, throughout the years and despite my geographical mobility, that of vehicular networking. The latter is a term that I will use extensively in this manuscript, but that, quite unfortunately, has been used with diverse meanings in the literature. Therefore, I would first like to disambiguate it: throughout this document, the expression *vehicular networking* is used as a proxy for *all realizations of networking systems that involve some form of communication from, to, or between vehicles*. In this acceptation, vehicular networking is not limited to the popular vehicular ad hoc network (VANET) paradigm, where vehicles build a spontaneous network by leveraging direct vehicle-to-vehicle communication and without need for fixed infrastructures. Instead, the definition of vehicular networking I will adopt encompasses, among others: (i) hybrid VANETs where vehicles can also communicate with fixed access points typically referred to as roadside units (RSUs); (ii) delay-tolerant networks (DTN) where vehicles participate in the store-carry-and-forward communication; (iii) systems that implicate direct communication between vehicles and the cellular radio access network (RAN) infrastructure – nowadays increasingly referred to as the networked vehicle scenario.

My interests in this broad connotation of the vehicular networking concept are rooted in my doctoral studies. Thus, I believe that a coherent presentation of the framework of my research cannot but begin by outlining the work I conducted during my PhD – which I will do in Sec. 1.2. Also, considering that I first undertook PhD studies back in 2005, it is quite natural that the research context has matured significantly since that moment: Sec. 1.3 provides an outlook on how vehicular networking has evolved from a fundamental research topic to an (almost) commercial paradigm. How my postdoctoral activities fit in such an emerging context is instead the focus of Sec. 1.4, where I will introduce the different subjects detailed in the remainder of the document. The organization of the latter is finally summarized in Sec. 1.5.

1.2 Early research topics

One day in January 2005, when I was at the very beginning of my PhD, my advisor at Politecnico di Torino, Prof. Claudio Casetti, invited me to have a look at the emerging research field of vehicular ad hoc networks (VANETs), i.e., completely autonomous wireless networks composed of communication-enabled vehicles. These were envisioned to be equipped with one or more radio interfaces, whose physical (PHY) and Medium Access Control (MAC) layer operations would be inherited from popular WiFi standards, possibly over dedicated frequency bands [2, 3]. The main use case for VANETs was the implementation of Intelligent Transportation Systems (ITS) solutions aiming at improving road safety and reducing traffic congestion. Indeed, specific consortia were emerging at that time in the US [4], in Europe [5], and in Japan [6], and academic-minded⁴ researchers were already starting to propose much more complex VANET applications, services, and overlays.

⁴No positive or negative nuance meant.

A fortuitous participation to the second edition of the ACM Workshop on Vehicular Inter-NETworking, Systems, and Applications⁵ allowed me to discover this lively research environment, dispelling my doubts⁶ and pushing me to plunge headfirst into this new, fascinating subject. As a result, between 2005 and mid 2008, I worked on several aspects of vehicular networking. The original and main subject of my PhD concerned the design of protocols for data sharing within VANETs, as discussed in Sec. 1.2.1. Since simulation was⁷ the only mean to evaluate such protocols, I also developed tools to simulate VANETs – namely, their underlying node mobility – in a more realistic way, as detailed in Sec. 1.2.2. In occasion of a long-term visit abroad, I also had a chance to experiment with real-world communication-enabled vehicles: the work presented in Sec. 1.2.3 allowed me to test in the field the actual performance of the WiFi-based technologies that I assumed in my different simulation-based studies.

1.2.1 Information management in VANETs

Under the supervision of Prof. Casetti and with the support of Prof. Carla-Fabiana Chiasserini, I mainly worked on the management of contents in VANETs. We developed distributed solutions for the retrieval, storage and dissemination of data in vehicular ad hoc networks, either pure or hybrid. More precisely, we developed an application framework based on a content-centric approach, named InfoShare [7, 8], allowing the exchange of small-sized data among vehicles. Then, we implemented on top of InfoShare two protocols, Eureka [9, 10] and Hamlet [11, 12], which leveraged estimates of the geographical presence of content in order to guide the propagation of queries and to make decisions on caching information at each node, thus making the data exchange more effective at no additional communication cost.

1.2.2 Mobility in vehicular networks

As my experience with vehicular networking technologies deepened, I started to grow the feeling that using random waypoint⁸ was probably not the best way to model the movement of vehicles when simulating my application-level data management protocols. This consideration led me to learn the basics of vehicular traffic flow theory, with a focus on microscopic mobility modeling, and to exploit such knowledge to code, jointly with Jérôme Härri (at that time a PhD student at Eurecom, France, and now an Associate Professor at the same institution) a generator of more realistic vehicular mobility traces for network simulation, named VanetMobiSim [13, 14, 15]. Although I had originally planned to use VanetMobiSim for my own research⁹, the

⁵Better known as the ACM VANET workshop, and soon to become the reference meeting of a dedicate and very active international research community.

⁶Partially due to the unfathomable relationship between VANETs and multipath transport protocols in wired networks, i.e., my original assignment as a research associate waiting to start of his doctoral studies. Yet I will never thank Claudio enough for that shift of direction.

⁷And still is nowadays.

⁸Random waypoint is a stochastic mobility model that describes the movement of each entity as an alternation of constant-speed linear movements and stationary pause periods. Although hardly representative of any real-world movement pattern, it was widely used in studies on generic mobile ad hoc networks (MANETs) – of which VANETs were considered a subset.

⁹Which I did, of course.

tool met significant success within the research community, and became one of the most popular mobility generators for vehicular network studies [16].

Getting acquainted with the world of vehicular traffic flow theory led Jérôme and myself to grow the belief that unrealistic mobility represented an extremely dangerous building ground for the research on vehicular networking in general, and on VANETs in particular. Therefore, we started asking ourselves if it were possible to quantify the impact of the level of realism of the vehicular mobility representation on the performance evaluation of VANET solutions. We ultimately opted for a protocol-agnostic analysis of the relationships between road traffic dynamics and the topological properties of the resulting VANET. We could show that the simplistic representations of mobility that were commonly employed in network simulation could heavily bias the connectivity of the VANET, affecting in turn the performance of protocols in unpredictable ways [17, 18].

1.2.3 Vehicular handover in urban-scale mesh networks

During my PhD studies, I was lucky enough to spend eight months at Rice University, in Houston, TX, USA, working with Prof. Edward Knightly and Tasos Giannoulis. There, we conducted experimental research on the support that large-scale wireless mesh networks¹⁰ could provide to users onboard vehicles and to networked vehicles. We first carried out a wardriving-style campaign to measure the coverage of the urban mesh network that Rice University had deployed in an under-resourced¹¹ neighborhood of Houston. Then, we designed, implemented and evaluated client-side-only WiFi handoff schemes that significantly improved the data transfer rates experienced by vehicular users accessing the Internet through the mesh network [19].

1.3 Evolution of the research context

After graduation, a significant part of my research activities continued along the directions outlined by the work carried out during my PhD thesis. However, the context had significantly evolved in the meanwhile: as it is natural for a young and increasingly popular scientific research field, that evolution had implied an expansion of the existing research directions, the opening of novel ones, and a slow but steady approaching of less futuristic but more practical network engineering solutions. In fact, such an evolution is still ongoing nowadays, and, despite important advances and the appearance of the first commercial applications, the full maturity of the field seems to be still to come.

More in detail, over the last decade, a mix of unreliable performance evaluations, disappointing experimental results, and pressure from the telecommunication

¹⁰Also referred to as WMNs, these networks are composed of fixed WiFi relays connected to gateways via a wireless multi-hop backbone. Such relays can then bring Internet connectivity to mobile clients that are far from the actual gateways in a cost-efficient and flexible way, without the need to deploy wired landlines.

¹¹I later discovered that the neighborhood was considered under-resourced because of the average yearly income of around \$20,000 of its residents. I prefer not to comment on how my PhD student salary compared to that amount.

and automotive industry has led vehicular networking studies to progressively shift away from the pure VANET paradigm¹² that initially drove the efforts on the subject. Attention has moved towards communication-enabled vehicles that employ at a time multiple wireless technologies, and that make large use of the existing as well as of forthcoming telecommunication infrastructures. The original principle of a spontaneous, distributed, self-organizing vehicular network is still there, but it is now considered as a complement to a pervasive communication from, to, and among vehicles via the cellular RAN. Indeed, the former can provide short-range, fast, inexpensive, opportunistic data transfers that integrate the seamless coverage of the strongly centralized cellular architecture. As a consequence, the large attention on scarcely practical problems, such as that of routing unicast data in a VANET [20], has left the way to, e.g., experimental assessments of the performance of vehicle-to-vehicle and vehicle-to-infrastructure communications [21, 22].

Application use cases have similarly matured. In the context of ITS, high-performance traffic management solutions have appeared that completely rely on cellular communications only, no matter if they are based on traditional commercial [23, 24] or on more original crowdsourcing [25, 26] approaches. The adoption of other networking technologies for traffic management services seems thus improbable. On the other hand, centralized cellular technologies, including third-generation Universal Mobile Telecommunications System (UMTS) and fourth-generation Long-Term Evolution (LTE), have shown non-negligible limits in terms of delay and capacity when safety-oriented services, such as cooperative awareness and danger notification, are considered [27, 28]. There, significant space is available for vehicle-to-vehicle communication solutions, which may in fact be merged onto future fifth-generation access network standards.

In the context of non-safety applications, the evolution has been even more evident. Visionary proposals for Internet-inspired peer-to-peer content exchange overlays built on VANETs have been abandoned in favor of more sensible use cases, such as the dissemination of small, disruption-tolerant contents to a large number of users onboard cars [29], the offloading of the cellular infrastructure via vehicle-to-vehicle data exchanges [30], and the gathering, fusion and exploitation of floating car data¹³ concerning the vehicles themselves as well as their surrounding environment [31], which makes vehicles active participants of the global urban sensing architecture.

The inclination to pragmatism has also brought to the creation of new international consortia dedicated to vehicular networking, such as NoW [32] in Europe and the continuing ASV in Japan. Specific frequency bands have been allocated to direct vehicular communication in these regions, as well as in the US. Also, standardization bodies, such as IEEE, IETF, OSI and ETSI have drafted new standards especially designed for vehicular Direct Short-Range Communication (DSRC), such as IEEE 802.11p [33], IEEE 1609 [34], OSI CALM-M5 [35] and ETSI ITS [36]. More recently, the automotive industry has become interested in the subject as well, as testified by the involvement of all major national car manufacturers in the deployment of

¹²Which remains nonetheless rather popular within the academic research community.

¹³Also known as FCD, a term indicating generic data originated at vehicles, which is then typically uploaded to Internet-based processing centers for knowledge discovery.

the largest live vehicular network testbed ever, within the framework of the German Sim^{TD} project [37].

Overall, all these initiatives suggest that we are approaching the moment when the first vehicular networking technologies will finally hit the mass market.

1.4 Selected research results

The research activities on vehicular networking I carried out during the past six years originate from those I initiated during my PhD, in Sec. 1.2 above. In fact, they consist in more functional and reasonable networking approaches to problems that represent themselves evolutions – oriented at more practical use cases – of those I addressed during my PhD. In that sense, I believe that my work has been coherent with the progress of the whole research field discussed in Sec. 1.3.

Next, I will briefly overview these research activities. A more detailed technical discussion of the same is provided in Chapters 2 to 4. Please note that, below as well as in the aforementioned Chapters, I classify the results of my research according three major categories. That choice was made for the sake of readability, and I hope it does not becloud the fact that there exist a continuity throughout my research.

1.4.1 Scaling up mobility modeling and assessing its impact

Characterizing the mobility of vehicles and understanding how it affects vehicular network solutions is, in my opinion, one of the fundamental problems of vehicular networking. As a matter of fact, road traffic dynamics determine the positions of vehicles at each time instant, which in turn draw the vehicular network topology. And, as taught in undergraduate classes on computer networking, networks with different topologies may yield very diverse strength and weaknesses, which communication protocols must exploit or compensate for, respectively. It is thus paramount that we comprehend which are the salient topological features of a network built on vehicles, and design and evaluate dedicated protocols in a way that accounts for such features.

Within such a context, my works on the simulation of vehicular mobility for network purposes, sketched in Sec. 1.2.2, paved the road to several studies I carried out with two PhD students I co-advised at INSA Lyon, i.e., Sandesh Uppoor¹⁴ and Diala Naboulsi¹⁵. Specifically, the awareness of the risks linked to the use of unrealistic mobility models pushed us to design and generate synthetic road traffic datasets that yield realistic features from all points of view. Once publicly released, such datasets would have finally allowed the research community to run reliable simulative evaluations of vehicular network protocols and architectures.

That rationale led us to the development of a large-scale synthetic dataset of vehicular mobility, describing 24 hours of road traffic in a 400-km² region comprising the conurbation of Cologne, Germany. The dataset, built using realistic information

¹⁴Sandesh, co-advised with Prof. Fabrice Valois, successfully defended his PhD thesis in November 2013, and is presently a postdoctoral fellow at Orange Labs, in Paris, France.

¹⁵Diala, co-advised with Dr. Razvan Stanica and Dr. Hervé Rivano, is currently at her second year of doctoral studies, and her defense is expected by late 2015.

on the road topology and macroscopic travel demand, and including over 700,000 individual trips, significantly advanced the state of the art. A 2-hour sample¹⁶ was made freely available via the project website¹⁷.

The availability of such a dataset allowed us to investigate, at a much larger scale than previously possible, the impact of road traffic dynamics on vehicular networking technologies. Specifically, we unveiled fundamental properties of both ad hoc and infrastructure-based vehicular communication systems. In the first case, we described the structure of the instantaneous topology of a vehicular network, as well as its spatiotemporal variations. In turn, this allowed to comment on the availability and reliability of pure ad hoc vehicular networks in our reference scenario. In the second case, we characterized the pervasive access from networked vehicles to the current cellular infrastructure deployed in the region. We assessed geographical and temporal network-wide access features, such as the offered load distribution and its movement patterns, as well as microscopic cell-level properties, such as inter-arrival and residence time distributions.

1.4.2 From VANET traffic to data management for networked cars

The availability of reliable mobility datasets – and the improved understanding of the properties they induce on the vehicular network – allow the design of more sensible network solutions, as well as their evaluation in dependable settings. Within that context, I targeted the most basic service provided by vehicular network solutions, i.e., transferring data towards and from vehicles.

Data management represents indeed a major focus of my activities, as a post-doc first and as a faculty at INSA Lyon later on. Building on past experience on information management in VANETs I addressed the problem of how to enable efficient data transfers between communication-enabled vehicles and different types of access network infrastructures. As a matter of fact, my studies on pure VANETs made me aware that, despite all efforts, the ad hoc network paradigm imposed too strong limitations on the nature and volume of data exchanges among vehicles: it could hardly provide latency or delivery guarantees, or support transfers over long distances, or manage but small-sized¹⁸ contents. Thus, I started considering systems where communication-enabled vehicles could also rely on some infrastructure, employing vehicle-to-vehicle communication as a complement to either a dedicated architecture composed of DSRC-based RSUs or the existing UMTS/LTE cellular network.

A revision of the data traffic model was also in order, since unicast data exchanges among vehicles did not appear to map to any real-world service. Instead, I studied two data traffic models of practical interest. The first is that of the download of

¹⁶The limitation on the dataset portion we could publicly distributed was imposed to us by the research partner who provided key source information used in the generation process.

¹⁷<http://kolntrace.project.citi-lab.fr/>.

¹⁸By small-sized, I intend here contents that do not exceed 100 KB. Although these may not seem small in some contexts (e.g., Internet browsing, or distributed sensing) they become so when considering bandwidth-hungry applications such as video streaming – which, by the way, dominate mobile data traffic [38].

large-sized contents from Internet-based servers to the vehicles, which is motivated by the realistic expectation that passengers on networked cars or buses will behave exactly as home-based network users, and access resource-intensive applications from their vehicles. The second is that of the upload of small-sized but high-frequency data from the vehicles to Internet-based processing centers, which represents the data traffic scenario that a popular mobile crowdsensing service would induce.

Jointly with Prof. Casetti, and Prof. Chiasserini of Politecnico di Torino, Italy, and with Dr. Francesco Malandrino, a PhD student and then a postdoctoral fellow at the same institution, we investigated the performance limits of data downloading to vehicular users, leveraging optimal combinations of infrastructure-to-vehicle and vehicle-to-vehicle communications. Practical protocols for such a cooperative download process were instead proposed in collaboration with Dr. Oscar Trullols-Cruces and Prof. Josè Maria Barcelo-Ordinas of Universitat Politècnica de Catalunya, Spain. Concerning the upload scenario, a joint work with Dr. Razvan Stanica (first as a postdoc at ENSEEIHT, France, and then as a colleague at INSA Lyon) brought us to explore the bounds of the improvement that vehicle-to-vehicle communication can bring to the upload of Floating Car Data (FCD). Within the same activity, we also proposed distributed algorithms that approximate such bounds.

1.4.3 The need for secure positioning

While studying data management protocols for vehicular networks, one common requirement emerged as especially critical for the efficient operation of the schemes we devised: the knowledge of vehicles' locations. In fact, not only data management, but a whole range of network solutions for vehicular environments heavily rely on information about the position of vehicles. A few representative examples include ITS techniques for danger warning or traffic monitoring, geocasting protocols, and techniques for accident liability attribution.

The need for positioning information implicitly mandates that vehicles must share their current location; failure in doing so can have a dramatic negative impact on protocol operation. Even worse, injection of false positioning information may easily disrupt the whole vehicular network, or let adversarial users gain illicit advantages. Therefore, practical vehicular networks require dependable positioning systems, where the location announced by vehicles is verified, and malfunctioning or misbehaving nodes are properly identified.

Motivated by the relevance of secure positioning to the vehicular networking context, I carried out studies on the verification of the location information advertised by nodes in vehicular networks in order to attain a so-called cooperative awareness, which is in turn the basis for multiple road safety services. Such studies were the result of collaborations with Dr. Francesco Malandrino, Prof. Chiasserini and Prof. Casetti of Politecnico di Torino, Italy, and with Prof. Panos Papadimitratos of KTH, Sweden. They addressed the definition of position verification protocols based on the concept of periodic beaconing, in presence of either fully distributed or centralized systems. In both scenarios, we developed algorithms that exploit physical properties of wireless communications in order to bound vehicles to their actual locations.

1.5 Document contents and structure

The remainder of the manuscript is organised according to the three research activity categories identified above. Specifically:

- Chapter 2 discusses our studies on the generation of realistic large-scale vehicular mobility datasets, and present the main results on the impact of road traffic on spontaneous and infrastructure vehicular network paradigms. This maps to the research sketched in Sec. 1.4.1 above. I opted for starting the technical part of the document from these topics because they deal with the very fundamental properties of vehicular networking, which are application-agnostic and thus of more general validity than the specific network solutions described in the following Chapters.
- Chapter 3 introduces our works on the management of data transfers briefly introduced in Sec. 1.4.2. The Chapter will address the problem in the downlink and uplink directions separately, since the two problems are intrinsically different.
- Chapter 4 provides some detail on the solutions we proposed for securing positioning information in vehicular networks, considering both distributed and centralized scenarios, as per Sec. 1.4.3.
- Chapter 5 presents very recent research topics that represent both continuity as well as a change in direction with respect to the vehicular networking studies that characterized the majority of my work over the past six years. Introducing these current research activities will also allow outlining my future research plans.
- Appendix A contains a complete curriculum vitae presenting my academic background, teaching and advising experiences, research projects and a full list of publications.

As a final remark, I would like to stress that I wrote this document trying to keep a balance between readability and technical depth. That was not a trivial task, mainly because of the diverse nature of the addressed topics – some of which were clear enough without delving into the details, some requiring a thorough discussion, and some others simply too complex to allow for a complete presentation. I hope that the reader will bear with me if some solutions proposed in the next Chapters appear just outlined while others are covered more in detail: that was done purposely, with the aim of maintaining the document enjoyable and sound at the same time.

Bibliography

- [1] Inria UrbaNet – Réseaux capillaires urbains, Activity Report 2013, <http://raweb.inria.fr/rapportsactivite/RA2013/urbanet/>.
- [2] IEEE 802.11 Task Group p - Wireless Access for the Vehicular Environment (WAVE), http://grouper.ieee.org/groups/802/11/Reports/tgp_update.htm.

- [3] IEEE 802.11 Task Group r - Fast Roaming/Fast BSS Transition, http://grouper.ieee.org/groups/802/11/Reports/tgr_update.htm.
- [4] Vehicle Infrastructure Integration (VII), <http://www.its.dot.gov/vii>.
- [5] Car-2-Car Communication Consortium (C2C-CC), <http://www.car-to-car.org>.
- [6] Advanced Safety Vehicle Program, Phase III (ASV-3), <http://www.mlit.go.jp/road/ITS>.
- [7] M. Fiore, C. Casetti, C.-F. Chiasserini, "On-demand Content Delivery in Vehicular Wireless Networks," IEEE/ACM MSWIM, Montreal, Canada, October 2005.
- [8] M. Fiore, C. Casetti, C.-F. Chiasserini, M. Garetto, "Analysis and Simulation of a Content Delivery Application for Vehicular Wireless Networks," Elsevier Performance Evaluation, 64(5), June 2007.
- [9] M. Fiore, C. Casetti, C.-F. Chiasserini, "Efficient Retrieval of User Contents in MANETs," IEEE INFOCOM, Anchorage, AK, USA, May 2007.
- [10] M. Fiore, C. Casetti, C.-F. Chiasserini, "Information Density Estimation for Content Retrieval in MANETs," IEEE Transactions on Mobile Computing, 8(3), March 2009.
- [11] M. Fiore, F. Mininni, C. Casetti, C.-F. Chiasserini, "To Cache or Not To Cache?," IEEE INFOCOM, Rio de Janeiro, Brazil, April 2009.
- [12] M. Fiore, C. Casetti, C.-F. Chiasserini, "Caching Strategies Based on Information Density Estimation in Wireless Ad Hoc Networks," IEEE Transactions on Vehicular Technology, 60(15), June 2011.
- [13] J. Härri, M. Fiore, F. Filali, C. Bonnet, "VanetMobiSim: Generating Realistic Mobility Patterns for VANETs," ACM VANET, Los Angeles, CA, USA, September 2006.
- [14] M. Fiore, J. Härri, Fethi Filali, Christian Bonnet, "Vehicular Mobility Simulation for VANETs," SCS/IEEE ANSS, Norfolk, VA, USA, March 2007.
- [15] J. Härri, M. Fiore, F. Filali, C. Bonnet, "Vehicular Mobility Simulation with VanetMobiSim," Transactions of The Society for Modeling and Simulation, 87(4), April 2011.
- [16] S. Joerer, C. Sommer, F. Dressler, "Toward Reproducibility and Comparability of IVC Simulation Studies: A Literature Survey," IEEE Communications Magazine, 50(10):82–88, October 2012.
- [17] M. Fiore, J. Härri, F. Filali, C. Bonnet, "Understanding Vehicular Mobility in Network Simulation," IEEE MoVeNet, Pisa, Italy, October 2007.

- [18] M. Fiore, J. Härrri, “*The Networking Shape of Vehicular Mobility*,” ACM MobiHoc, Hong Kong, China, May 2008.
- [19] A. Giannoulis, M. Fiore, E. Knightly, “*Supporting Vehicular Mobility in Urban Multi-hop Wireless Networks*,” ACM/USENIX MobiSys, Breckenridge, CO, USA, 2008.
- [20] F. Li, Y. Wang, “*Routing in vehicular ad hoc networks: A survey*,” IEEE Vehicular Technology Magazine, 2(2):12–22, June 2007.
- [21] D. Hadaller, S. Keshav, T. Brecht, S. Agarwal, “*Vehicular Opportunistic Communication Under the Microscope*,” ACM MobiSys, San Juan, Puerto Rico, USA, June 2007.
- [22] F. Bai, D.D. Stancil, H. Krishnan, “*Toward Understanding Characteristics of Dedicated Short Range Communications (DSRC) From a Perspective of Vehicular Network Engineers*,” ACM MobiCom, Chicago, IL, USA, September 2010.
- [23] TomTom, “*Travel Time Measurements using GSM and GPS Probe Data*,” White Paper, 2010.
- [24] Meihui Software, <http://www.meihuichina.com>.
- [25] Coyote Systems, <http://www.coyotesystems.co.uk>.
- [26] Waze – outsmarting traffic together, <http://www.waze.com>.
- [27] T. Mangel, T. Kosch, H. Hartenstein, “*A Comparison of UMTS and LTE for Vehicular Safety Communication at Intersections*,” IEEE VNC, Jersey City, NJ, USA, December 2010.
- [28] G. Araniti, C. Campolo, M. Condoluci, A. Iera, A. Molinaro, “*LTE for Vehicular Networking: A Survey*,” IEEE Communications Magazine, 51(5):148–157, May 2013.
- [29] C. Lochert, B. Scheuermann, M. Mauve, “*Information Dissemination in VANETs*,” in H. Hartenstein and K. Laberteaux (editors), VANETs: Vehicular Applications and Inter-Networking Technologies, pp. 4980, John Wiley Sons, 2010.
- [30] J. Whitbeck, Y. Lopez, J. Leguay, V. Conan, M. Dias de Amorim, “*Relieving the wireless infrastructure: when opportunistic networks meet guaranteed delays*,” IEEE WoWMoM, Lucca, Italy, June 2011.
- [31] M. Fiore, A. Nordio, C.-F. Chiasserini, “*Investigating the Accuracy of Mobile Urban Sensing*,” IEEE WONS, Banff, AL, Canada, March 2013.
- [32] Network-on-Wheels, Federal Ministry of Education and Research, Germany.
- [33] IEEE 802.11p published standard. “*Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular Environments*.” July 2010.

- [34] IEEE 1609. “*Family of Standards for Wireless Access in Vehicular Environments (WAVE)*.” U.S. Department of Transportation, January 2006.
- [35] OSI 21215 Standard. “*Communications access for land mobiles, (ITS-CALM-M5)*.”
- [36] ETSI Standard. “*EN 302 665 Intelligent Transportation Systems (ITS)*.”
- [37] Sim^{TD}, <http://www.simtd.de>.
- [38] Cisco, “*Global Mobile Data Traffic Forecast Update, 20122017*,” White Paper, February 2013.

Chapter 2

Fundamental properties of vehicular networking

We define as *fundamental properties* those characteristics that are inherent to the nature of the vehicular network, i.e., that are independent of the different technologies, protocols, or applications employed. Such properties are unique to vehicular networks, since they mainly stem from the dynamics of road traffic: indeed, the mobility of vehicles is hardly comparable to that of other mobile network nodes. Since they directly determine the basic capabilities of the communication system, understanding the fundamental properties of vehicular networks is paramount to the sensible design of dedicated protocols and architectures.

The results presented in this Chapter concern the generation of realistic large-scale vehicular mobility datasets, in Sec. 2.1, and its exploitation in order to assess the impact of road traffic on spontaneous and infrastructure vehicular network paradigms. The former are presented in Sec. 2.2, where we characterize the instantaneous topology of a vehicular network, while the latter are introduced in Sec. 2.3, where we discuss the properties of pervasive vehicular access to the cellular infrastructure.

These works were the outcome of collaborations with Sandesh Uppoor, Diala Naboulsi, Oscar Trullols-Cruces, Josè Maria Barcelo-Ordinas, and Jèrôme Härri. They resulted in the following publications:

- S. Uppoor, M. Fiore, “*Large-scale Urban Vehicular Mobility for Networking Research*,” IEEE VNC, Amsterdam, The Netherlands, November 2011.
- S. Uppoor, M. Fiore, *Insights on metropolitan-scale vehicular mobility from a networking perspective*,” ACM HotPlanet, Low Wood Bay, UK, June 2012.
- S. Uppoor, M. Fiore, J. Härri, “*Synthetic mobility traces for vehicular networking*,” Vehicular Networks: Models and Algorithms, H. Labiod and A.-L. Beylot (Editors), Wiley, 2013.
- D. Naboulsi, M. Fiore, “*On the Instantaneous Topology of a Large-scale Urban Vehicular Network: the Cologne case*,” ACM MobiHoc, Bangalore, India, July 2013.

- S. Uppoor, O. Trullols-Cruces, M. Fiore, J.M. Barcelo-Ordinas, “*Generation and Analysis of a Large-scale Urban Vehicular Mobility Dataset*,” IEEE Transactions on Mobile Computing, to appear.

2.1 Modeling road traffic for vehicular networking

Vehicular networking solutions investigated by the research community include algorithms, protocols and architectures that cover the whole network stack. They address, e.g., RSU deployment, medium access and transmission power control, data rate adaptation, client-driven handover management, admission control, multi-hop and disruption-tolerant routing, reliable and best-effort data traffic transport, and file sharing. Applications have also been proposed to target the different usages related to road safety and traffic management as well as, e.g., broadcasting, content downloading, data dissemination, urban sensing, or cellular offloading. The performance evaluation of most such solutions requires large-scale scenarios, making direct experimental assessments impractical due to their cost and complexity. Simulation becomes then the tool of choice in the validation of new network architectures and protocols for vehicular environments.

2.1.1 State of the art

Unfortunately, simulative performance evaluation of vehicular networks are often biased by the underlying mobility representation. As a matter of fact, as repeatedly proven in the past [1, 2, 3], the movement of vehicles can dramatically affect the behavior of network protocols, and an incorrect representation of car traffic can lead to misleading conclusions, even in presence of a flawless network-level simulation. As a result, it is today acknowledged that, for the results of a vehicular simulative campaign to be credible, mobility traces must be employed that capture the unique macroscopic and microscopic dynamics of car movement patterns.

Such considerations have led to substantial progress in the quality of car movement traces for vehicular networking research over the last few years. The simplistic stochastic models employed in early works [1] have been replaced by random mobility over realistic road topologies [4] at first, and by microscopic vehicular models borrowed from transportation research [5] later on. These features were then included in dedicated simulation environments, and integrated with road signalization [6, 7]. Ever since, vehicular mobility simulators have been growing their complexity and features, allowing to accurately simulate the individual movement of vehicles over realistic road topologies [8]. Moreover, in parallel with the evolution of synthetic traces of vehicular mobility, real-world car traffic dataset have grown in number and scale.

Tab. 2.1 summarizes the most relevant features of the different mobility traces currently employed for network simulation. We categorize the datasets based on the nature of their macroscopic traffic data, i.e., the sources employed to determine the time and routes of trips traveled by individual vehicles in the dataset. None of such datasets satisfies the requirements imposed by a reliable vehicular network

Table 2.1: Major features of the vehicular mobility traces currently available for network simulation

Synthetic trace	Simulator	Traffic data	Granularity (s)	Duration (h)	Area (km ²)	Roads	Vehicles	Availability
Porto [9]	DIVERT	perception	1	0.3	62	major/minor	5000	upon request
Turin [10]	SUMO	perception	1	1	20	major/minor	1500	upon request
Karlsruhe [11]	VanetMobiSim	perception	1	3	21	major/minor	2000	upon request
Zurich [12]	GMSF	perception	1	0.5	12	major/minor	420	online
Berlin [13]	Videlio	measurements	1	14	21.56	major/minor	955	upon request
Porto [14]	-	photography	snapshot	-	41.3	all	10566	upon request
Berkeley [16]	-	detectors	60	72	-	single freeway	-	upon request
Toronto [16]	-	detectors	20	24	-	single freeway	-	upon request
Bologna [17]	SUMO	detectors	1	1	20.6	major/minor	10333	within consortium
Luxembourg [18]	SUMO	detectors	1	12	1700	major	150000	online
Portland [19]	TRANSIMS	survey	1	0.25	21	major/minor	16529	not available
Braunschweig [20]	VISUM/VISSIM	survey	1	6	12	major/minor	5000	not available
Canton of Zurich [21]	MMTS	survey	1	6+6	65000	major	260000	online
Real-world trace		Traffic data		Duration (days)		Type		
Seattle [23]	-	AVL/GPS	120	13	5100	bus	1200	online
Chicago [24]	-	AVL/GPS	40	17	606	bus	1648	online
DieselNet [25]	-	GPS	1	60	241.40	bus	30	online
ShanghaiGrid [26]	-	GPS	60	1	102	taxi	1171	online
ShanghaiGrid [27]	-	GPS	n.d.	4	150	bus	700	upon request
Beijing [28]	-	GPS	60	1	750	taxi	2927	upon request
T-Drive [29]	-	GPS	180	7	750	taxi	10357	online
San Francisco [30]	-	GPS	60	30	18000	taxi	500	online
Frankfurt [31]	-	GPS	-	-	300	cars	420	not available

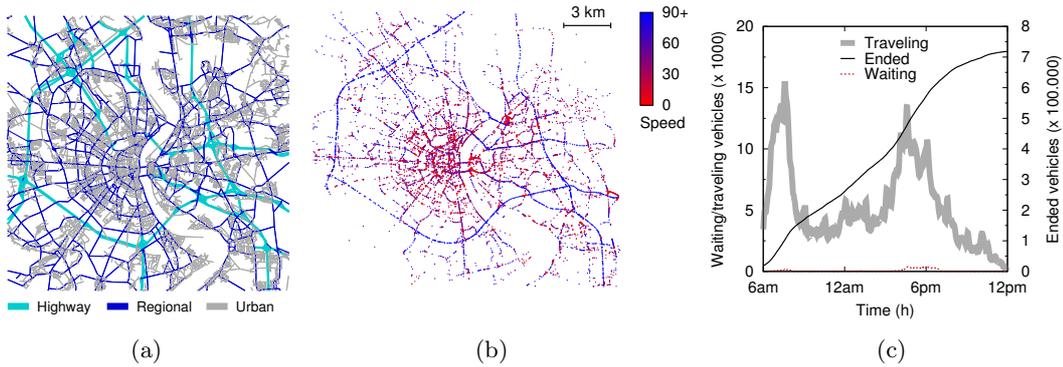


Figure 2.1: Cologne mobility dataset features. (a) Road network layout. (b) Road traffic snapshot at 7:00 am. (c) Simulated vehicles status over time.

simulation. Specifically, these traces are only representative of single highway environments [16], just focus on small regions and short timespans [9, 10, 11, 12, 14, 13, 17, 19, 20], rely on incomplete travel demands [18], employ simplistic microscopic modeling [21], or only cover a subset of the road traffic with low time granularity [23, 24, 25, 26, 27, 28, 29, 30, 37].

2.1.2 The Cologne dataset

We thus went a step further, generating an original synthetic dataset of urban vehicular mobility, that is characterized by an unprecedented combination of scale, detail and realism. The dataset covers the city and surroundings of Cologne, a typical medium-sized European city located in Germany, and was obtained by coupling the well-known state-of-art tools detailed next.

- First, realistic maps of the road topology in the region are extracted from the OpenStreetMap (OSM) database [32]. The OSM map includes information on major road infrastructures such as roads of different types, railways, buildings, but also data on road signalization, including traffic lights and stop or yield signs. All this information is contributed by a vast community through satellite imagery and GPS traces, and is commonly regarded as the highest-quality road data publicly available today. The OSM data employed in the Cologne dataset covers an area of 400 km^2 around the urban agglomeration, including around 4500 km of roads. The complex road network of the Cologne metropolitan areas is depicted in the left plot of Fig. 2.1(a), where highways, regional and urban roads are told apart through different colors and line thickness.
- Second, microscopic-level realism is granted by a faithful representation of the behavior of individual drivers, in terms of acceleration and braking in presence of other cars nearby, while respecting at the same time the road regulations, such as speed limits and right-of-way's. Such low-level mobility is simulated via the Simulation of Urban Mobility (SUMO) software [33], an open-source traffic simulator developed by the German Aerospace Center (DLR). SUMO leverages

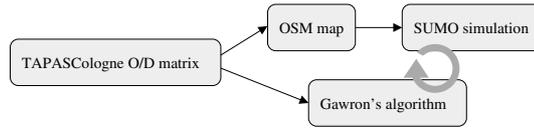


Figure 2.2: Cologne mobility dataset generation workflow.

Krauss’ car-following model [34] and Krajzewicz’s lane-changing model [35] to regulate a driver’s acceleration and overtaking decisions. These models, long validated by the transportation research community, make of SUMO one of the most scalable and complete open-source microscopic vehicular mobility generators available today.

- Third, the Cologne dataset is realistic also from a macroscopic point of view, i.e., it faithfully reproduces the large traffic flows across the Cologne metropolitan region. To that end, the dataset leverages a realistic *traffic demand*, represented by a so-called Origin/Destination (O/D) matrix, i.e., a list of the start time, origin and destination of each car trip in the simulated area. In order to generate such an O/D matrix, real-world data collected in the Cologne region by the German Federal Statistical Office – including 30,700 daily activity reports from more than 7000 households [36, 37] – is processed through the Travel and Activity PAtterns Simulation (TAPAS) methodology [38]. Finally, Gawron’s *traffic assignment* model [39] is run on the O/D matrix, so as to identify the actual route followed by each driver.

The components presented above are combined as depicted in Fig. 2.2 in order to generate the vehicular mobility dataset. The information contained in the TAPAS O/D matrix are initially used to identify the boundaries of the exact simulation region, extract the associated map from OSM and filter it so as to remove unneeded content that does not concern the road layout. Then the OSM map is converted to a format readable by SUMO, and fed to the microscopic mobility simulator. The TAPAS O/D matrix is also used as an input to Gawron’s algorithm, which, in turn, determines an initial traffic assignment and provides it to SUMO. Then, a first vehicular mobility simulation can be run with SUMO, and, once finished, a feedback on the resulting traffic density over the road topology is sent back to Gawron’s algorithm. Based on such new information, a new traffic assignment is computed, and a second SUMO simulation is run. The process is repeated until a stable traffic assignment is generated that allows to accommodate the whole volume of the traffic demand¹.

The resulting dataset comprises more than 700,000 car trips in the Cologne conurbation over a period of 24 hours, mimicking the road traffic activity during a typical working day. Fig. 2.1(b) displays the road traffic recorded in the dataset at 7:00 am. We can note that highways appear in bright blue, corresponding to speeds higher

¹We remark that the dataset generation required in fact a very time-consuming process of integrating the different tools and fixing a number of inaccuracies, errors and incompatibilities. For more details, please refer to our paper presented at IEEE VNC 2011.

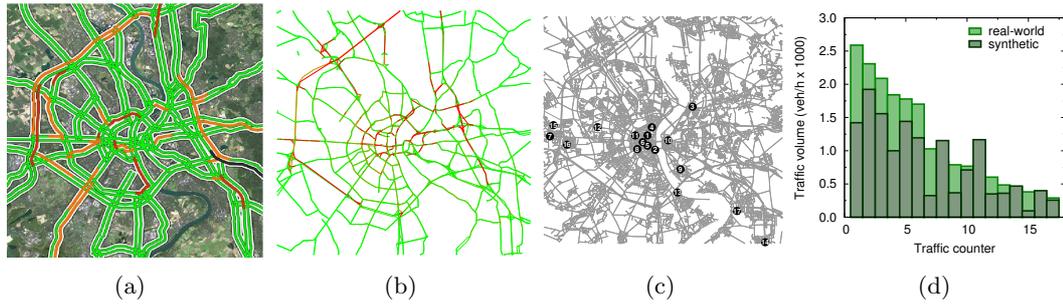


Figure 2.3: Cologne mobility dataset validation. (a) Real-world traffic at 5:00 pm. (b) Synthetic road traffic at 5:00 pm. (c) Measurement locations of road traffic. (d) Daily traffic volumes at each such location.

than 90 km/h. Large portions of the urban roads are in violet, indicating fluid traffic conditions. Congestion appears only in the city center, where dark red regions are visible that map to speeds from 30 to 50 km/h.

The macroscopic properties of the dataset are shown in Fig. 2.1(c). We can remark that the number of traveling cars follows typical daily activity patterns, with peaks during the morning (from 7:00 am to 9:00 am) and afternoon (from 4:30 pm to 6:00 pm) rush hours. An approximate maximum of 15,000 vehicles travel at the same time over the road topology, at around 8:00 am. Real-world behaviors, such as very low traffic at night and a lower traffic peak at around noon, can also be observed. Also, the number of ended trips grows over time, as more and more drivers reach their destinations, and the number of vehicles waiting to enter the simulation is reduced to values close to zero. Both are proofs of the correctness of the simulation, since the road network can correctly accommodate the travel demand.

Interestingly, we found the macroscopic traffic simulated in the Cologne dataset to nicely match that observed in the real world, through live traffic information services. In Fig. 2.3(a) and Fig. 2.3(b), we compare the road traffic information retrieved through the ViaMichelin service at 5:00 pm with the simulation output at the same hour. This represents a critical period of the day, in the middle of the afternoon traffic peak, and key features of real-world mobility patterns are faithfully reproduced in the dataset: e.g., the congestion on the highways around the city, where commuters merge with long-distance travelers passing through the region, or generalized but discontinuous heavy traffic in the city center, especially along major roads.

We also retrieved traffic counts from several traffic hotspot locations in Cologne [40]. The data was collected by the local transportation authority, and refers to a typical working day, thus matching the situation our synthetic scenario refers to. The measurement locations are displayed in Fig. 2.3(c), whereas the actual traffic counts – expressed in vehicles per hour, and averaged over the whole day to match the format of real-world counts – are shown in Fig. 2.3(d). There, for each location, we report two bars, the lighter referring to the real-world data and the darker to the simulated data. We can observe a rather good match between the microscopic behavior of the synthetic dataset and that recorded in Cologne: absolute values are

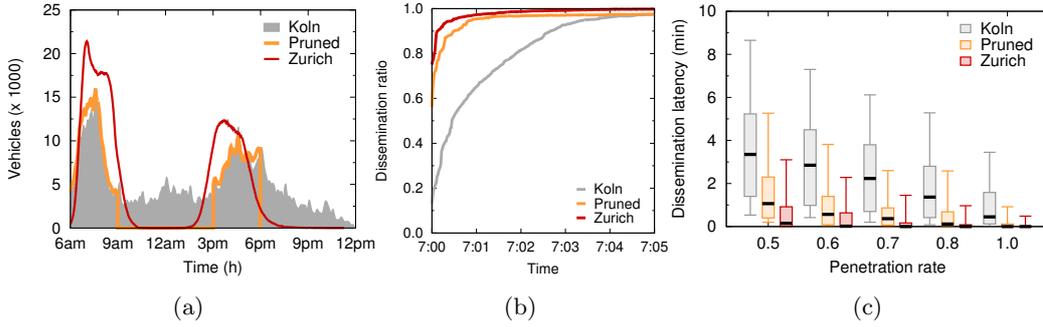


Figure 2.4: Comparative evaluation of the vehicular network performance in presence of different mobility datasets. (a) Road traffic over time. (b) Epidemic dissemination ratio over time with a penetration rate of 1.0. (c) Latency to achieve the dissemination ratio quantiles 0.05, 0.25, 0.5, 0.75 and 0.95 under different penetration rates.

close enough, and relative ones generally respect the traffic volume ordering of the real-world measurements.

2.1.3 Relevance of road traffic realism

An interesting question is whether the additional realism brought by the Cologne dataset has an impact on the evaluation of vehicular network solutions, and, if so, to which extent. In this Section, we consider one specific application use case, i.e., the epidemic dissemination of a small content² in the urban area through vehicular opportunistic communication, and observe the effect that different mobility representations have on the system performance. We compare the Cologne dataset against two mobility traces:

- The Zurich trace, is a 400-km² subset of the Canton of Zurich dataset introduced in [21]. This is a very popular road traffic trace that has been widely employed by the networking research community. However, it has several limitations. First, the road map is coarse, only accounting for highways and main traffic arteries in Zurich. Second, the Multi-agent Microscopic Traffic Simulator (MMTS) used for the trace generation employs a queuing approach [22], faster but less accurate than the car-following model we adopt. Third, the traffic demand is not as accurate as the one we dispose of, as portrayed in Fig. 2.4(a). There, we compare the traffic volume recorded in our dataset with that in the Zurich region scenario. Although the general trend is the same, with traffic peaks early in the morning and in the middle of the afternoon, the traffic demand in the Zurich region scenario unrealistically drops to zero between 10:00 am and 2:00 pm, as well as after 8:00 pm. Clearly, the Zurich region dataset models the traffic peak hours only.

²This could map to, e.g., data on the current status of road traffic in the region or a map update for the on-board navigation system.

- The Pruned trace is a simplified version of our Cologne dataset. It is generated using the same OSM road topology, TAPAS travel demand and traffic assignment introduced in Sec. 2.1.2. However, it only accounts for major road arteries, and adopts a simplified microscopic model comparable to that used by MMTS. Finally, the travel demand is limited to the morning and afternoon traffic peaks, as displayed in Fig. 2.4(a). The rationale behind the Pruned trace is to have a dataset that yields similar characteristics to the Zurich one in the same urban area of our Cologne dataset.

In all mobility scenarios, the epidemic dissemination is started by a source located at the city center, which broadcasts the content for the first time at 7:00 am, i.e., during the morning rush hour. We test different technology penetration rates, i.e., ratios of cars equipped with vehicle-to-vehicle communication interfaces and participating in the network.

Fig. 2.4(b) portrays the dissemination ratio, i.e., the percentage of vehicles reached by the content, versus time, when all the vehicles take part in the dissemination process. We observe that the dissemination is very fast in all scenarios, as almost all vehicles are informed in a few minutes. However, while the curve obtained through the Cologne dataset is more gentle, with two minutes required to reach 80% of the network, the Zurich and Pruned traces lead to a much faster spreading. In fact, 80% of the vehicles are informed in a few seconds in these scenarios.

The performance difference holds also at different penetration rates: Fig. 2.4(c) summarizes the initial spreading performance of the dissemination, by reporting the latency in reaching different quantiles of the dissemination ratio. For each penetration rate, in abscissa, we compare the time required to reach 5%, 25%, 50%, 75% and 95% of the vehicles, in each mobility scenario. The content is successfully disseminated throughout the whole network in all cases, however it is also clear that the dissemination is much faster in the case of the Zurich and Pruned traces, with latencies that are two to six times lower than those recorded in the Cologne scenario.

We conclude that:

- the additional realism granted by the Cologne dataset has a significant impact on the performance of the vehicular network in the application scenario considered. In fact, we claim that this result is of general validity: further investigations, omitted here for the sake of brevity, showed that the reason behind such performances is that the underlying mobility greatly impacts the connectivity of the vehicular network. The latter in turn affects any service running on top of it. Therefore one cannot dependably rely on popular yet simplistic datasets of vehicular mobility for networking studies.
- the Zurich and Pruned mobility traces result in very comparable outcomes. This proves that it is not the underlying urban environment that determines the differences between the Cologne and Zurich scenarios, rather the diverse level of realism of the mobility description.

2.2 Instantaneous topology of a vehicular network

The observation that vehicular mobility can impact dramatically the topology of vehicular networks, and thus the outcome of performance evaluations carried out on networking solutions, led us to investigate more in detail the interactions between mobility and ad hoc connectivity in vehicular environments. In particular, we focused on the characterization of the major topological features of vehicular networks. Our aim was to answer to key questions that, despite some previous efforts [41, 9, 42, 43], remained open, such as: *is the vehicular network well connected or highly partitioned? Which size can clusters of multi-hop connected vehicles attain? Which is the internal structure of such clusters? How sparse or dense are single-hop communication neighborhoods? How do all these network connectivity features vary in time? How do they depend on the geographical location?*

The responses to these questions directly determine the strengths, weaknesses and overall capabilities of a spontaneous vehicular network, and shall thus be among the main drives to the design of dedicated protocols. Moreover, they are the key to quantifying the *availability* and *reliability* of the network, i.e., the main concerns of car manufacturers when it comes to vehicle-to-vehicle multi-hop communication. In order to attain our goal, we modeled the vehicular network by sampling the Cologne mobility dataset at every 10 seconds, and mapping each road traffic snapshot to one instantaneous connectivity graph, where vehicles are the nodes and communication links among them are the edges of the graph³. The resulting set of graphs thus provides a complete description of the geographical and temporal dynamics of the vehicular network topology over 24 hours in the Cologne conurbation. We then borrowed tools from complex network theory that have been successfully employed to characterize a number of large-scale real-world networks such as those of Internet routers, World Wide Web pages or interacting social species [44, 45, 46], and applied them to our set of graphs. Next, I will discuss the main results of our analysis, by separating network-, component-, and node-level observations.

2.2.1 Network level

The level of connectivity of the whole vehicular network is mainly characterized through two metrics: the number \mathcal{C} of components⁴, that is an index of the level of network fragmentation, and the component size \mathcal{S} , which describes the heterogeneity of the fragmentation.

The Cumulative Distribution Function (CDF) of \mathcal{C} , aggregated over all the vehicular network graphs extracted from the whole 24 hours of road traffic, is portrayed in Fig. 2.5(a). In 80% of cases the vehicular network has more than 1,000 disconnected components, and the inset plot shows that most of the probability is concentrated

³For a fully detailed, rigorous discussion of the instantaneous graph model, including considerations on the radio-frequency signal propagation modeling, we invite the reader to refer to our paper presented at ACM MobiHoc 2013.

⁴A component is defined as a set of nodes (i.e., vehicles) such that a (multi-hop) path exists between each two nodes in the set. Therefore, vehicle-to-vehicle communication is possible within each component, but different components are isolated from each other from a network viewpoint.

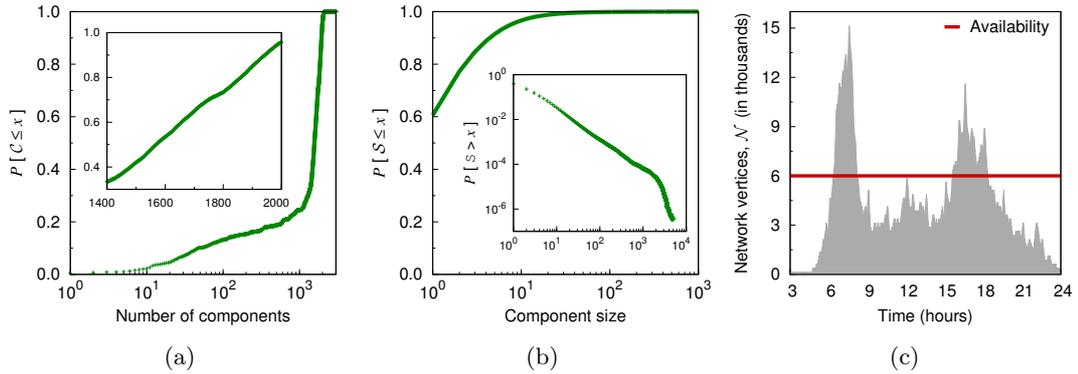


Figure 2.5: Network-level metrics. (a) Aggregate distribution of the number of components, \mathcal{C} . (b) Aggregate distribution of the size of components, \mathcal{S} . (c) Network availability over the day.

in a linear growth between 1,000 and 2,000 components. This suggests that the vehicular network is highly partitioned into thousands of separate node groups unable to communicate with each other.

The distributions of the component size \mathcal{S} help us clarify whether we face many components of similar size or a heterogeneous network of both large and small components. The CDF in Fig. 2.5(b) shows that the network is largely made of very small components, with 60% of them being *singletons*, i.e., isolated vehicles, and 95% of the components comprising 10 vehicles or less. However, by looking at the Complementary CDF (CCDF), portrayed in the inset plot, we can appreciate the heavy tail of the distribution, appearing as linear on a log-log scale. There is thus a non-negligible probability that the aforementioned small components coexist with components that include up to 2,000 vehicles connected through direct links or multi-hop routes. After such a component size, the CCDF has an exponential decay, and components as large as 4,500 nodes appear with significant lower probability.

Such a general view of a partitioned and heterogeneous network is however aggregated over time and space. In order to unveil the impact of the daytime and the differences between geographical areas, we show in Fig. 2.6 the instantaneous vehicular network fragmentation measured in the Cologne region at different hours. In each plot, every circle corresponds to one component, its diameter and color mapping to the size of the component it represents: broader circles map to larger components, while the color code is reported on the bars at the right of the plots. The resulting images give a rough, yet intuitive, idea of the behavior of the network connectivity evolution: early in the morning, i.e., before 6:00 am, the network is very partitioned and only small components of 40 vehicles or less are present. The morning traffic peak, between 7:00 am and 8:00 am, has a very positive impact on the network topology, with the appearance of very large components of thousands of vehicles and a diffuse presence medium-sized components of several tens of cars. That effect disappears later on, and large components do not reappear until the afternoon traffic peak, between 16:00 and 18:00, although a slightly increased connectivity is observed

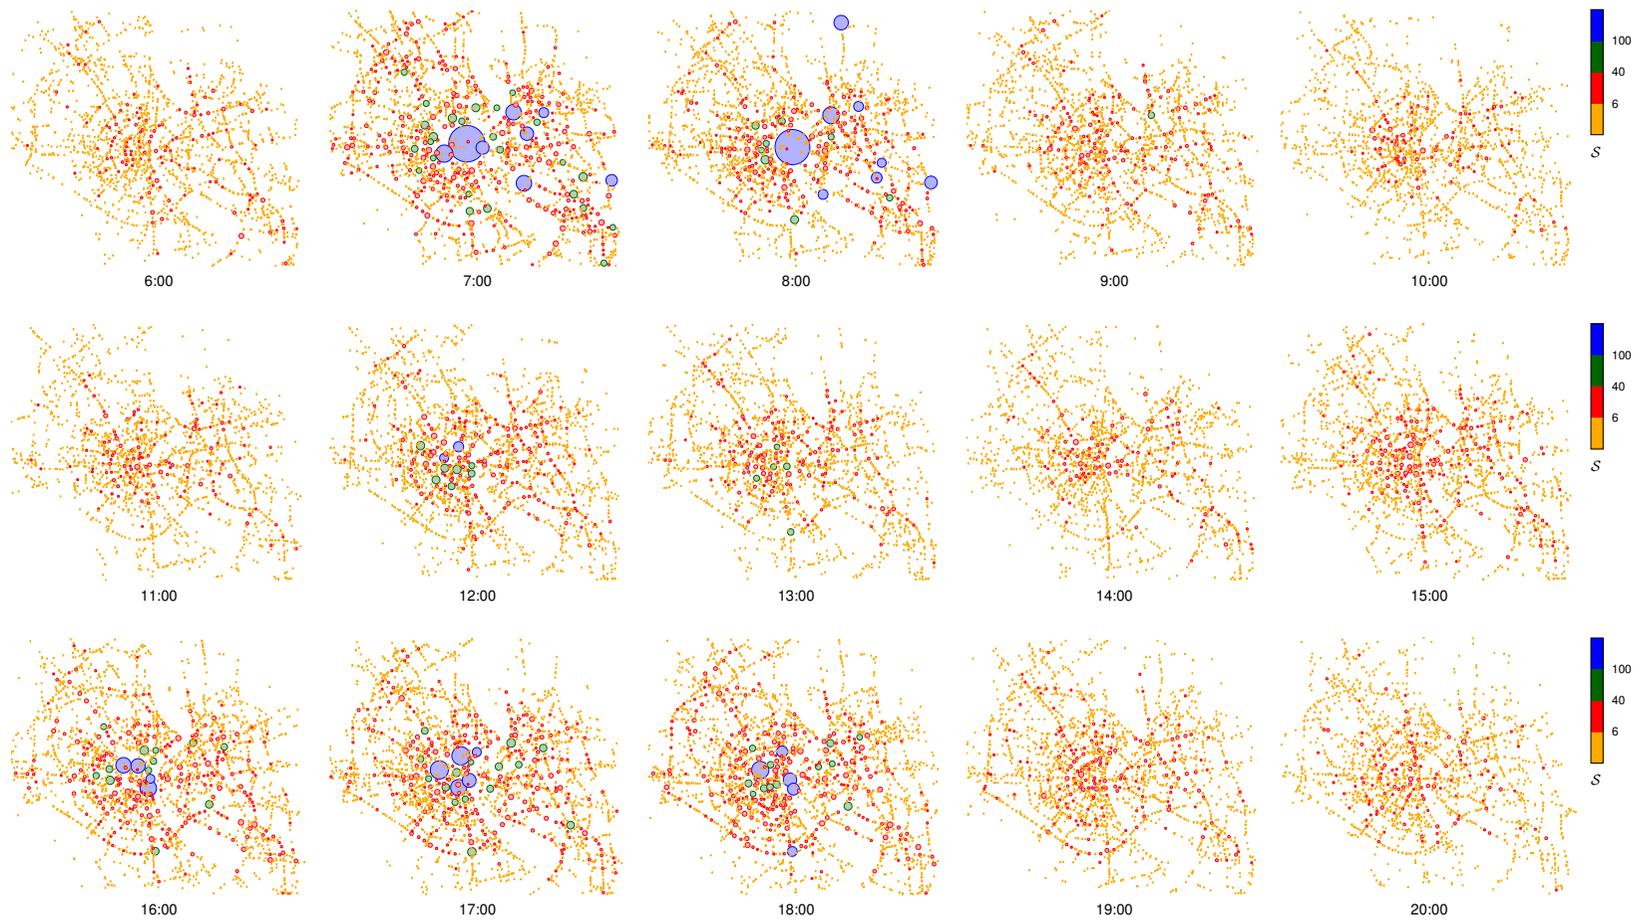


Figure 2.6: Spatiotemporal view of the network-level connectivity. Each plot is a geographical representation of the network components at a specific time of the day. This figure is best viewed in colors.

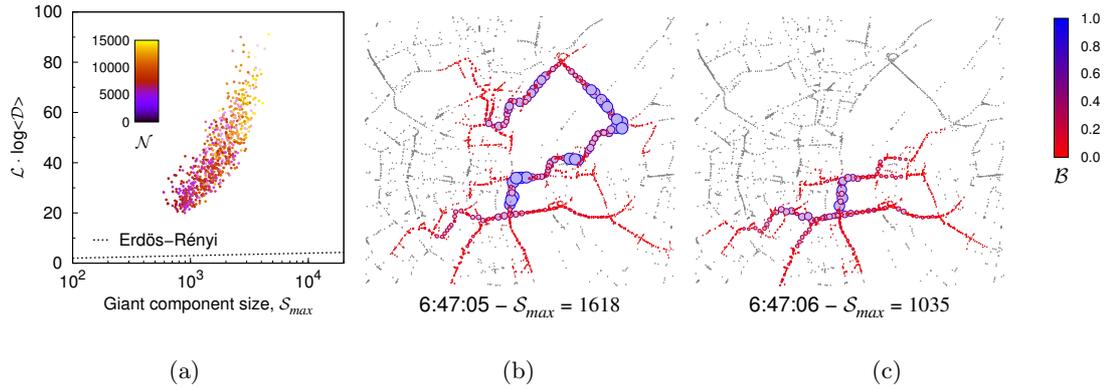


Figure 2.7: Component-level metrics. (a) Small world property of the vehicular network versus the number of graph nodes, \mathcal{N} . (b,c) Internal structure of the largest component at two subsequent time seconds, where vertices color and size indicate the node’s betweenness centrality.

around noon. Moreover, large components mostly appear in the city center, where the traffic is denser. We can infer that both time and space are paramount factors in the characterization of the vehicular network topology, which is generally highly fragmented, with larger components only appearing at specific locations during the traffic peak hours.

More formally, we observed that the network fragmentation is entirely driven by the number of nodes \mathcal{N} present in the graph, i.e., by the vehicular density in the urban region. In turn, this allowed us to individuate critical vehicular density thresholds above which the network is *available*, i.e., large clusters emerge that allow vehicle-to-vehicle multi-hop communication among hundreds of cars. This threshold is indicated as a solid line in Fig. 2.5(c), where it overlaps to the evolution of \mathcal{N} over time. Large clusters can be expected to appear in the vehicular network mainly during the morning and afternoon rush hours, with an availability of around 4 hours/day.

2.2.2 Component level

Not only the availability, but also the *reliability* of large components is a critical factor for the design of vehicular network solutions. In order to analyze this second aspect, we studied the one-second variability of the largest component size, S_{max} , and observed that it undergoes significant variations at each second, its value doubling or halving every few seconds. The reason behind these major size changes lies in the internal dynamics of large components. Although a rigorous discussion is not possible due to space limitations, we provide an intuitive example of such dynamics in Fig. 2.7(b) and Fig. 2.7(c). The two plots show the vertices in the studied large component at two subsequent seconds, and larger dot size and colors towards

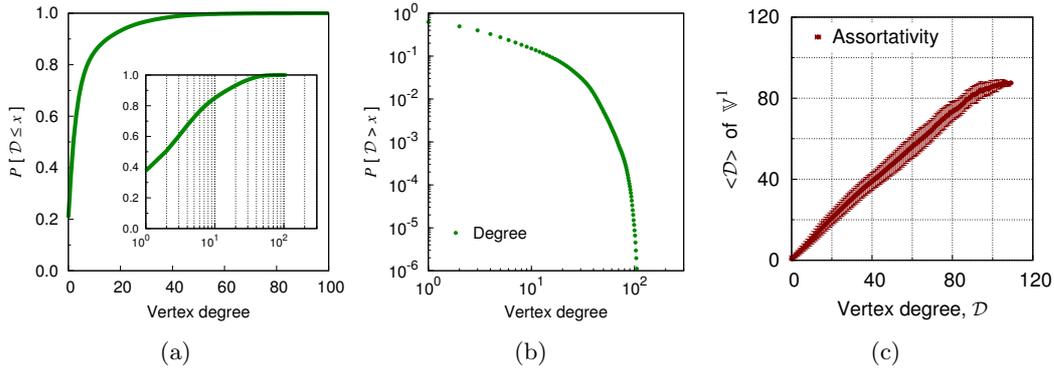


Figure 2.8: Node-level metrics, aggregated over the whole day. (a,b) CDF and CCDF of the vertex degree \mathcal{D} . (c) Vertex degree assortativity.

blue indicate higher betweenness centrality⁵. We observe that the largest cluster in Fig. 2.7(b) comprises over 1,600 vehicles, with a very high betweenness centrality of nodes in the northeastern region (where a bridge joins the two sides of the Rhine river). This implies that vertices in that area branch together large groups of vehicles. Indeed, at the next time second, in Fig. 2.7(c), a shift of a few vehicles in that same region is sufficient to break the connectivity over the bridge, disconnecting the whole group of vehicles located in the northern area. As a result, \mathcal{S}_{max} suddenly drops to around 1,000 nodes. These dynamics are very frequent in the vehicular network, which let us conclude that large components, when available, suffer from low reliability.

We further explore the structure of large components by studying their internal *navigability*, i.e., how easy it is to route data through vehicle-to-vehicle multi-hop communication within them. Fig. 2.7(a) compares the average length \mathcal{L} of the shortest paths between any two vertices within large components of different size \mathcal{S}_{max} (on the abscissa) recorded during the 24 hours. The value of \mathcal{L} quickly grows as \mathcal{S}_{max} increases, and it is always much higher than that of the equivalent Erdős-Rényi random graph, i.e., a typical example of small world network where each pair of nodes is separated by a low number of hops [44]. This confirms that larger components are not denser and thus better connected, but only geographically wider, possibly as a result of the temporary merge of smaller components. In conclusion such components are also very difficult to navigate as they grow in size [44, 45, 46].

2.2.3 Node level

At the individual node level, Fig. 2.8(a) and Fig. 2.8(b) show the CDF and CCDF of the node degree \mathcal{D} in the vehicular network, when aggregating all samples over 24 hours. The plots point out the heterogeneity of one-hop neighborhoods: around 80% of them are small (i.e., $\mathcal{D} < 5$), yet they co-exist with vertices with 50-100 neighbors. Also, as the degree distribution does not follow a power law, we conclude

⁵Betweenness centrality is a complex network theory metrics used to denote the criticality of a node for the global connectivity of a component.

that vehicular networks are not *scale-free* in the degree distribution [44, 46]. This means that the maximum one-hop neighborhood size is constrained (in our case, by geographical restrictions on the vertices positions) to sizes that are small with respect to the network size \mathcal{N} . In other words, although 50 neighbors may seem a lot, such large neighborhoods they only appear for high \mathcal{N} 's and represent a mere 1% of the whole network – again a sign of poor navigability.

The characterization of two-hop neighborhoods can be performed by observing the *assortativity* of the network, i.e., the correlation between the degree \mathcal{D} of a node and the average degree of its one-hop neighbors. Fig. 2.8(c) shows that the vehicular network is highly assortative: a strong linear relationship exists between these two measures, with a low standard deviation. This points out that high-degree vertices are typically connected to high-degree vertices, and low-degree ones to other low-degree vertices.

The combination of lack of scale-free property and high assortativity is a proof that, unlike many other real-world networks, vehicular networks do not show a backbone of high-degree hub nodes interconnecting low-degree leaf nodes. Rather, the network is structured into weakly tied cliques of cars with similar degree.

2.2.4 Discussion

We can summarize the main results of our analysis of the vehicular network topology as follows:

- the vehicular network is severely and consistently partitioned, which suggests that connected multi-hop routing and data dissemination are just not feasible through the whole network;
- different geographical locations and daytimes imply very diverse connectivity properties;
- very large network components do exist in the vehicular network, although with much lower probability and smaller size than previously believed [41, 42];
- such very large components are affected by: (i) low availability, as they only appear at specific locations during a few hours each day; (ii) poor reliability, as they are affected by a continuous merge-and-split process; (iii) little navigability, as vertices are connected by long and meandering multi-hop paths;
- vehicular communication neighborhoods are heterogeneous, and they display high assortativity, resulting in a network of weakly tied cliques of cars.

Such observations lead in turn to a number of considerations on the design of an effective vehicular network solutions. Namely, the high network fragmentation and the low availability and reliability of large components let us conclude that pure multi-hop connected communication is not a sensible strategy to adopt in urban vehicular networks. Instead, *carry-and-forward techniques are mandatory* to reach all network nodes and highly recommended even to route data within large components. The peculiar internal structure of the latter let us argue that *roadside unit (RSU) deployment is critical to maintain large components connected*, and thus increase their reliability and availability.

Even assuming the presence of bridging RSUs, large vehicular network components remain difficult to navigate, as they are sparse non-small worlds. Vehicles in these components are in fact connected through long articulated paths, most of which pass through a small set of bridging nodes. We thus deem that *geographic information is indispensable to data routing and dissemination* within complex connected components. Moreover *non-negligible end-to-end latencies are probably unavoidable* in a vehicular network of sensible extension, due to the length of multi-hop paths. A more controversial conclusion would be that, given the adversity of the vehicular network topology, *vehicles should resort to the cellular infrastructure most of the time*, except for localized transfers (within a few hundreds meters from the transmission origin) or specific situations (e.g., delay-tolerant dissemination during the rush hours).

Also, our results show that vehicles can move in a few tens of seconds from quasi-isolation to being part of cliques of several tens of cars. Then, even for the localized applications the vehicular network topology seems to support the best, *medium access control, data rate adaptation and power control algorithms will play a more important role than expected*. Indeed, the rapidity of such MAC-layer techniques to adapt to the varying network conditions may make the difference between an efficient network and a useless one.

Last but not least, the common practice of evaluating vehicular network protocols in small-scale scenarios with arbitrary vehicular densities shall be regarded as harmful.

As a final remark, we stress that all the results we derived cannot be considered valid but for the Cologne scenario we studied. Yet, they may be representative of many other urban regions of similar nature and size (i.e., middle-sized European cities). As an example, works on urbanism research has shown that metropolitan layouts can be easily grouped in a limited number of categories showing very similar properties [47].

2.3 Pervasive vehicular access

The results we just presented, on the fundamental features of vehicular networks based on a pure ad hoc communication paradigm, highlighted, among others, the relevance of the cellular radio access network (RAN) as a key infrastructure to enable communication among vehicles.

In fact, we anticipate that vehicles will play in the near future a non-negligible role in the explosion of the load on the RAN. On the one hand, the seamless Internet connectivity offered by the 3G/WiMAX network and the capacity boost promised by 4G/LTE are lurking on-board passengers into acting as home-based users, consuming and generating large amounts of data – possibly with quality-of-service constraints. On the other hand, vehicles themselves are being equipped today with their own communication interfaces, mainly used to upload so-called Floating Car Data (FCD). FCD consists in content that is generated by cars in an autonomous fashion, i.e., independently of their passengers, and is typically transferred to processing centers via the cellular infrastructure. Practical usages of FCD include today

real-time road traffic monitoring [23], fleet management [24] and distant support for safety, diagnostic and anti-theft services (e.g., BMW Assist, Ford SYNC, General Motor OnStar, Toyota Safety Connect, and Mercedes-Benz mbrace, just to cite a few representative examples). Yet, many more FCD-based applications are expected to follow in the years to come, making networked cars a key component of upcoming smart cities.

These trends will transform a vast majority of the cars that populate today's road traffic into very active mobile users. Such a transition will lead us into the age of pervasive vehicular access to the cellular RAN, where cars become a main source and destination of a large amount of mobile data traffic. Pervasive vehicular access will represent a significant challenge to network operators, given the combination of large data volumes, high-frequency access, elevate speed and unique movement patterns of vehicular users. For instance, WiFi and femtocell offload – the primary strategy for operators to ease the mobile data pressure, and expected to relieve the cellular network from almost 50% of the data traffic by 2017 [48] – will not be a solution for vehicular access. As a matter of fact, the limited communication range of such technologies and the travel velocity of cars lead to exceedingly short-lived contacts, whose median duration is limited to a mere 13 seconds according to experimental measurements [49]. That, the lack of handover support and the poor scalability of random channel access make WiFi unpractical for vehicular users.

Within such a context, it is important to anticipate how pervasive vehicular access will affect the cellular network. Specifically, understanding the unique dynamics of car-generated data traffic is a mandatory first step toward the design of dedicated solutions to accommodate the demand generated by highly mobile vehicular users. Such a characterization is the most critical in urban regions, where the mobile data traffic demand is already overflowing the RAN capacity, and where pervasive vehicular access will be at its highest. We thus provided a first characterization of the features of pervasive vehicular access in an urban environment, considering the urban region of Cologne as our case study. To that end, we leverage the mobility dataset introduced in Sec. 2.1, and information on the real-world deployment of the 247 base stations that build up the cellular access infrastructure of one telecommunication operator in the Cologne region [50]. Our analysis focused on the interactions between the vehicular mobility and the radio access network, investigating the offered load dynamics induced on the cellular infrastructure by on-board users and networked cars. Our approach is thus protocol- and application-independent, and let us draw results of general validity on the fundamental properties of pervasive vehicular access.

In this Section, I will first present our results on both macroscopic properties of vehicular access, in terms of expected traffic volumes and large-scale inter-cell flows, as well as microscopic features, namely inter-arrival and residence times at RAN cells.

2.3.1 Access load volumes

The spatiotemporal evolution of the pervasive vehicular access load is portrayed in Fig. 2.9. Plots refer to different hours of the day, between 5:00 am and 11:00 pm. In

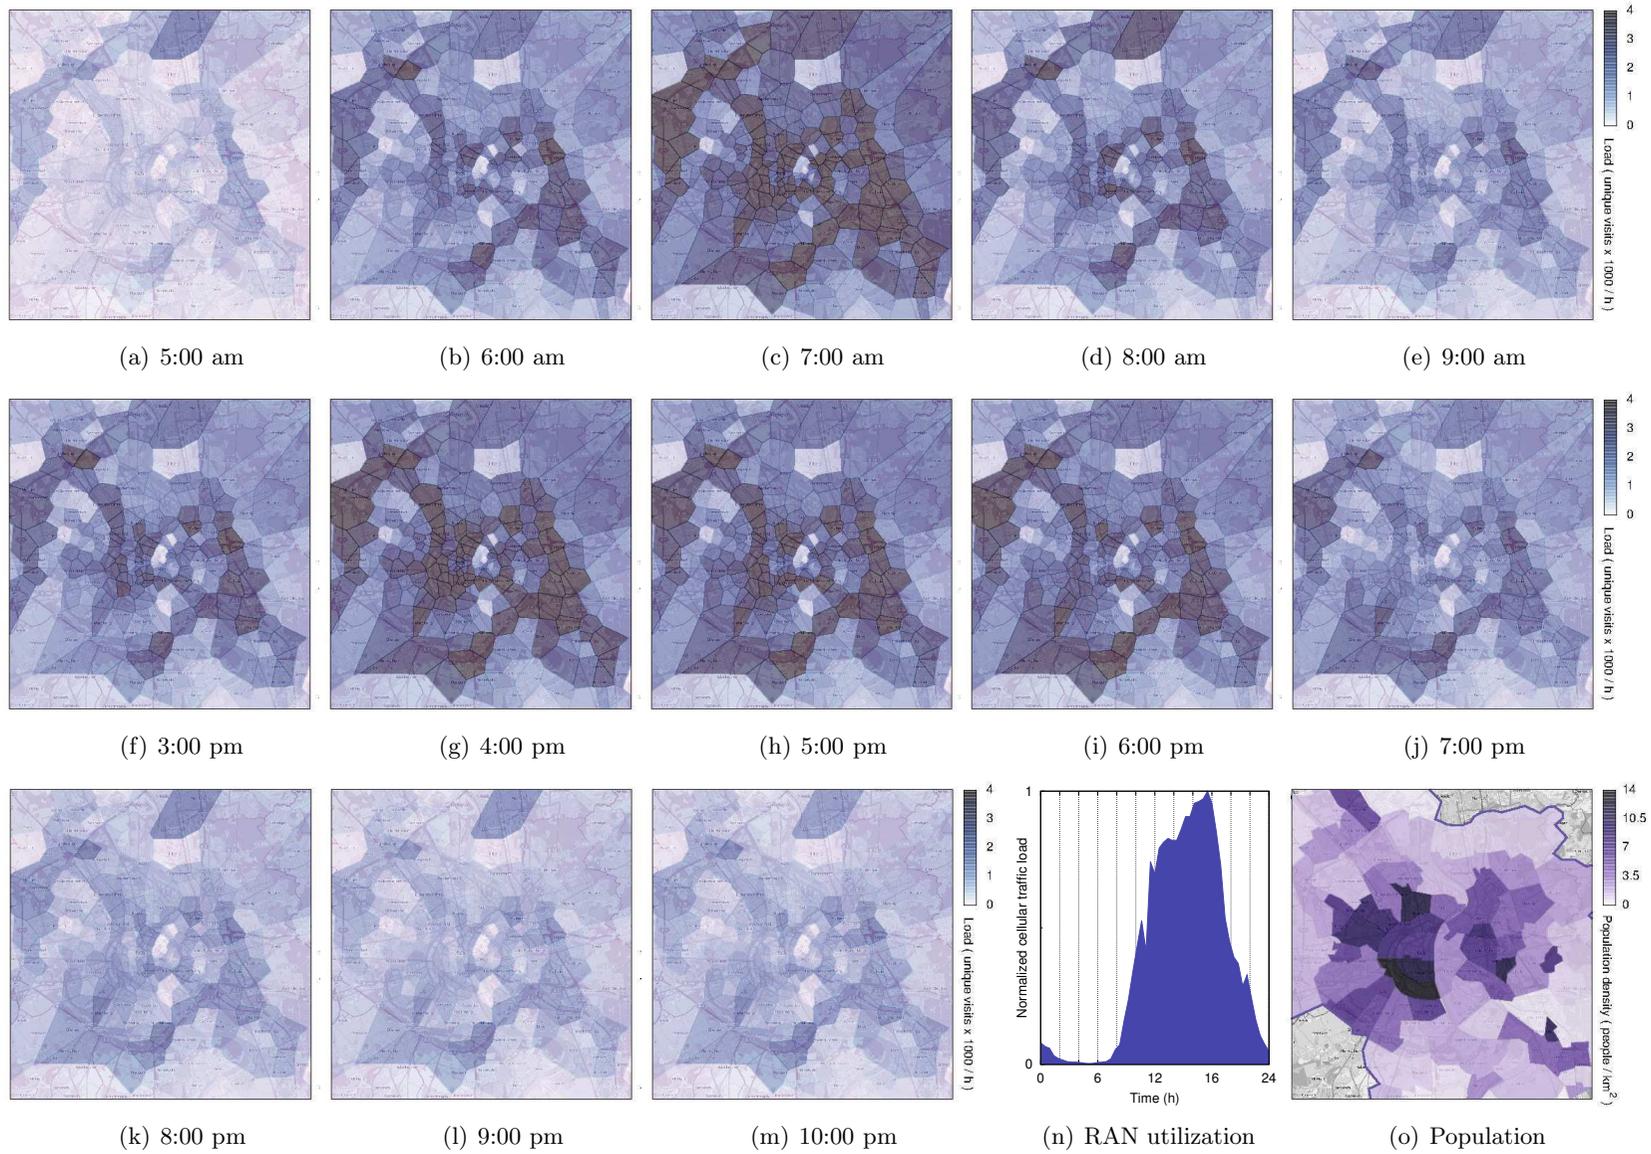


Figure 2.9: (a)–(m) Spatiotemporal evolution of the expected data traffic load generated by pervasive vehicular access in the Cologne region, during one day. (n) Typical daily profile of normalized RAN traffic load – concession of the Autonomous Networks Research Group at USC, <http://anrg.usc.edu>. (o) Geographical distribution of the population in the Cologne region. This figure is best viewed in colors.

each image, RAN cells are colored according to their associated load, according to the color range at the right of each row. We observe that very early in the morning the traffic is sparse throughout the whole region, with slightly higher density along the North-South highways that pass East and West of the urban area (see Fig. 2.1(a)). Road traffic grows rapidly from 6:00 am, and reaches a peak between 7:00 am and 8:00 am. The effect on the RAN is that base stations have to serve from several hundreds to several thousands additional users per hour. The precise value varies with the geographical location of the base station, and, as one would expect, strictly depends on the road network layout. Indeed, pervasive vehicular access load is the heaviest along the highways that surround the urban region as well as along primary thoroughfares that convey the road traffic towards the city center.

As we move past the rush hours, the vehicular access load is alleviated and then remains moderate during the morning, with a somehow higher access load is observed around midday⁶. The second daily peak is recorded during the afternoon, between 4:00 and 6:00 pm. Although less intense than the morning one, it proves to be much longer. Also, it is interesting to note that during the afternoon peak high loads are observed only along major roadways, while in the morning (Fig. 2.9(c)) the heavy traffic conditions force drivers to optimize their travel time by spreading more evenly over the street layout. The road traffic activity progressively decreases after 7:00 pm, and the vehicular access load with it.

We can remark the following features of interest from a RAN access viewpoint. First, the pervasive vehicular access shows a *strong daily variability*, with clearly identifiable peaks during the morning and afternoon. Although more intense, the morning peak is shorter and falls in a period of typical RAN under-utilization, as displayed in Fig. 2.9(n). Conversely *the afternoon access peak risks to be difficult to accommodate*, as it occurs when up to 95% of the RAN resources are already in use. Second, there is also an *elevate spatial heterogeneity* in the way vehicles will access the RAN, which is mainly driven by the road topology. If we assume that the usual access load can be mapped to the population density, in Fig. 2.9(o), we can observe that the geographical distributions of customary mobile traffic load and pervasive vehicular access do not overlap at all. This implies that *resource allocation in the RAN will need re-thinking* if vehicular users are to be served in a pervasive manner. On the positive side, capacity planning may be significantly eased by the relatively stable behavior of the geography of vehicle access. Indeed, the plots in Fig. 2.9 outline how the relative behavior of *low- and high-load areas remain the same* over the whole 24-hour period.

2.3.2 Access load flows

The previous results provided a static view of the spatial distribution of pervasive vehicular access at different daytimes. Considering the flows of road traffic grants a different perspective, unveiling the correlation of vehicular use mobility over the Cologne region. We thus now focus on macroscopic vehicular flows, i.e., large groups

⁶The plots referring to these hours are not shown due to page formatting constraints. The full picture is available within Sandesh Uppoor's PhD thesis, which is publicly downloadable at the following URL: tel.archives-ouvertes.fr/docs/00/91/25/21/PDF/thesis.pdf.

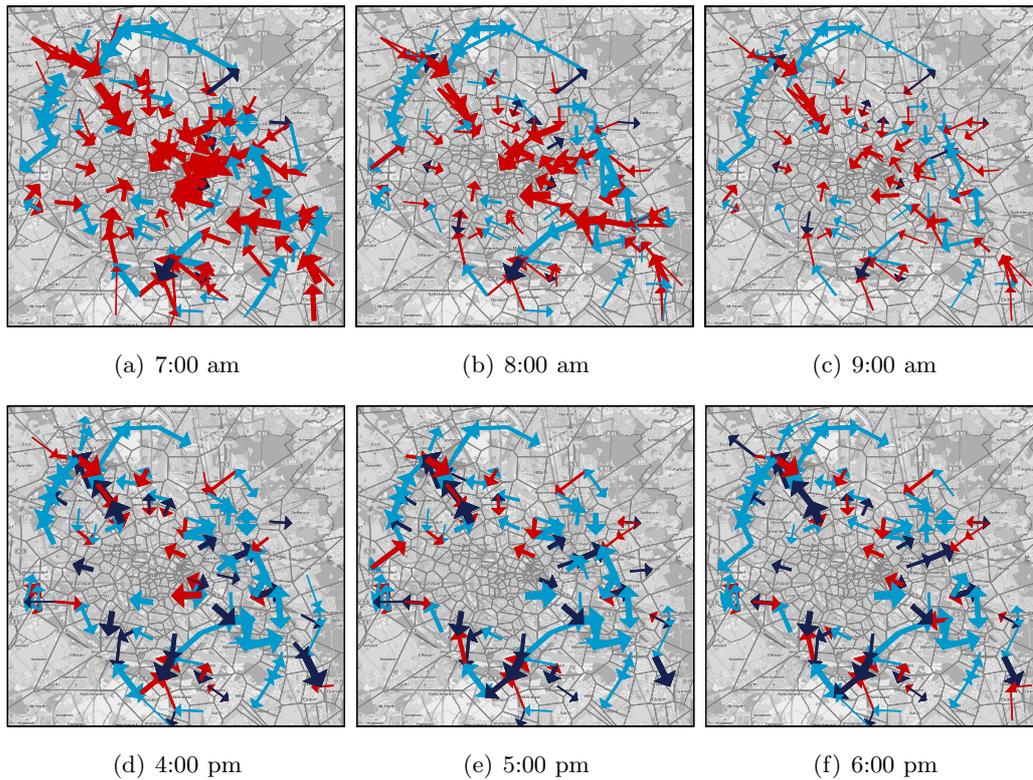


Figure 2.10: Spatiotemporal evolution of macroscopic flows of pervasive vehicular access in the Cologne region over one day. This figure is best viewed in colors.

of cars traveling along similar paths. More precisely, we focus on flows that occur between cells at a significant distance (i.e., more than one kilometer apart) and that have a non-negligible transition probability (i.e., more than 30% of the vehicular users move from the first cell to the second).

An intuitive geographical representation of such vehicular access flows is shown in Fig. 2.10, at different times of the day. In these plots, each arrow maps to one inter-cell flow. The arrow color is set depending on the flow direction: red arrows represent flows going towards the center of the Cologne urban area, dark blue arrows represent flows moving away from the center, and light blue arrows represent neutral flows that do not have a precise in- or out-bound direction. Finally, the size of the arrow is an indicator of the vehicular access volume associated to the flow, i.e., the number of vehicular users per hour following the flow: wider arrows map to larger flows.

Plots in Fig. 2.10(a)–(c) refer to the morning rush hours, between 7:00 am and 9:00 am. Major flows are detected along the primary roads that bring inbound traffic towards the city center, but also over the highways that surround the urban region. Instead, during the afternoon traffic peak, in Fig. 2.10(d)–(f), more heterogeneous traffic headings appear, and the largest significant flows are along the peripheral highways. Also, most of the traffic in the urban area moves this away from the

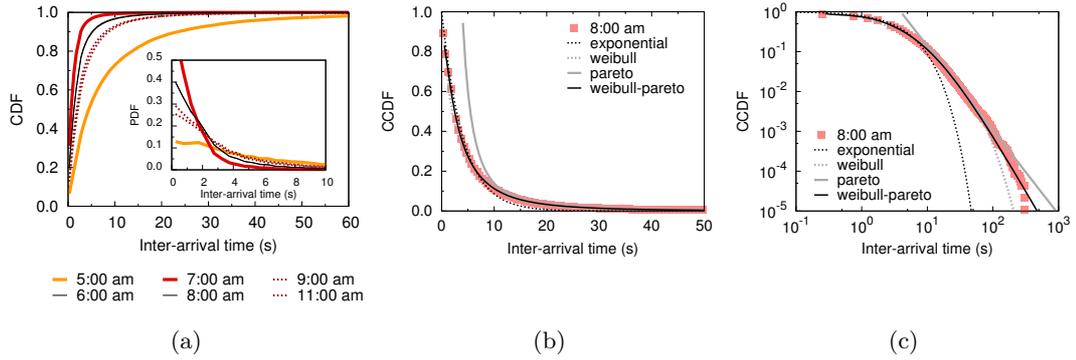


Figure 2.11: Distributions of vehicular user inter-arrival time at RAN cells. (a) Aggregate distributions over 30-minute intervals. (b,c) Theoretical fittings on the distribution recorded at 8:00 pm, in linear-linear and logarithmic-logarithmic scales.

city center, although the phenomenon is not as striking as that observed with the morning inbound flows. These behaviors are consistent with the daily activity cycle in all urban areas: in the morning the commuters entering the city merge with steady traffic over the highways, while the afternoon features flows leaving downtown.

From a RAN viewpoint, these results confirm the important spatiotemporal variability of the pervasive vehicular access. However, while the load density studied in Sec. 2.3.1 showed a stable geographical distribution of low- and high-load cells, vehicular user flows unveil how the movement that causes them is not stationary. Instead, *inter-cell mobility dynamics mostly change over daytime*, notably with three diverse behaviors during (i) the morning traffic peak, (ii) the afternoon rush hours, and (iii) the other periods of the day with low to moderate road traffic conditions⁷. Still, the stability periods of these major access flows last several hours each. Additionally some high-volume inter-cell flows that show day-long regularity can be identified at precise locations: this is the case, e.g., of freeways that surround the urban region. These observations let us conjecture that *macroscopic vehicular access flows can be leveraged towards the design of dedicated dynamic RAN resource allocation techniques*. Indeed, a proper understanding of the dynamics of pervasive vehicular access could turn mobility into an advantage, rather than an issue to cope with.

2.3.3 Cell inter-arrival times

Having characterized the macroscopic features of pervasive vehicular access, we now increase the detail level and focus on microscopic features that are of interest from a networking viewpoint. Specifically, we study vehicular access at the individual RAN cell level and the first property we consider is the vehicular user inter-arrival time. Jointly with the cell residence time, discussed in the Sec. 2.3.4, these metrics have been largely used in the past to characterize mobility in cellular networks.

We first evaluate the distribution of the inter-arrival time of vehicles at each

⁷Results not included, see footnote 6

RAN cell, in Fig. 2.11(a), during the morning hours⁸. We observe that inter-arrival time CDFs vary significantly at different hours of the day. The diversity is clearly correlated to the road traffic intensity. Low-traffic hours (e.g., 5:00) result in CDFs skewed towards higher inter-arrival times, while traffic-peak hours (e.g., 7:00 am) result in the lowest inter-arrival times. Similarly, the CDFs overlap for hours that yield comparable traffic conditions (e.g., from 9:00 to 11:00 am). In any case, inter-arrival times are typically very short, as expected of high-speed users. At least 95% of the vehicular users enter the coverage area of a cell within 10 seconds or less of each other. The interval drops to 3 seconds or less for 50% of the vehicular users.

One interesting observation is that, despite the diverse realizations, the CDFs and PDFs in Fig. 2.11(a) seem to share a common shape. We therefore study if a same theoretical distribution can fit the different empirical inter-arrival complementary CDFs (CCDFs) measured in the Cologne scenario. A representative sample of our results is provided in Fig. 2.11(b) and Fig. 2.11(c), at 8:00 pm, which corresponds to a mild road traffic scenario. The results confirm our intuition that the different empirical distributions follow similar laws, although with parameters that depend on the daytime considered. The linear-linear plot in Fig. 2.11(b) shows that a Weibull distribution fits very well the CCDFs for low inter-arrival times – noticeably better than an exponential one does. However, the empirical CCDFs display heavy tails, as proven by their linear shape in logarithmic-logarithmic plots. The Pareto class is thus the matching theoretical distribution for high inter-arrival times. Also in the tail region, the exponential distribution largely fails to reproduce the experimental behavior. Given these results, we employed a hybrid Weibull-Pareto distribution [51] to obtain a complete fitting over the whole inter-arrival time range. We found such a hybrid distribution to perfectly fit the empirical data at all times of the day, with a residual sum of squares (RSS) well below 0.1% for all the CCDFs.

Our results let us conclude that *inter-arrival times of vehicular users at RAN cells follow a hybrid Weibull-Pareto distribution* in the considered scenario. Although more tests in different urban environments are needed to grant general validity to the result, the latter is sufficient to invalidate the common assumption that user arrivals follow a Poisson process – and thus inter-arrival times are exponentially distributed – in the case of pervasive vehicular access. In addition, we observed how the *daytime induces significant differences in the inter-arrival time distribution*, and how pervasive vehicular access results in typical *average inter-arrivals in the order of a few seconds*.

2.3.4 Cell residence times

The cell residence time is the amount of time spent by a user under coverage of a same base station. We record the cell residence time of vehicular users in the Cologne scenario, separating different hours of the morning. The resulting distributions are reproduced in Fig. 2.12(a). Unlike what observed in the case of inter-arrival time, the cell residence time distributions measured at different hours basically overlap,

⁸The results for other hours of the day show similar properties and are omitted here for the sake of brevity.

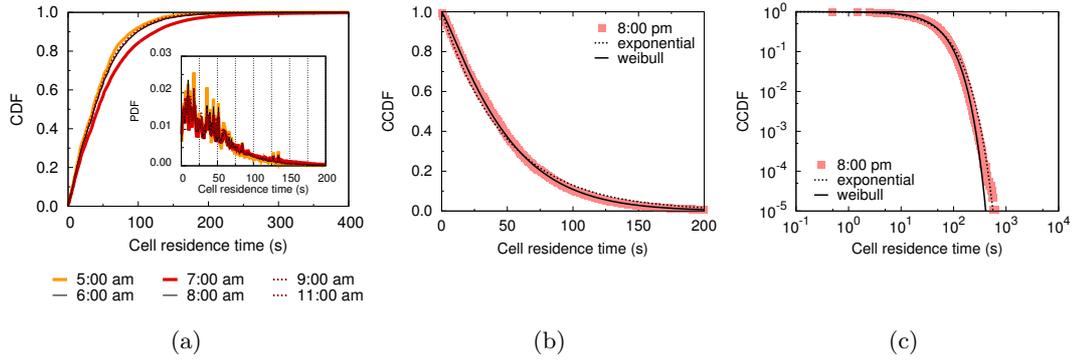


Figure 2.12: Distributions of vehicular user residence time within RAN cells. (a) Aggregate distributions over 30-minute intervals. (b,c) Theoretical fittings on the distribution recorded at 8:00 pm, in linear-linear and logarithmic-logarithmic scales.

i.e., the daytime does not seem to have a significant impact on the time spent by vehicular users in RAN cells. This may appear counter-intuitive when considering the important daily variations we observed in the access load and inter-arrival times. In fact, the result is just an indicator that the road infrastructure can accommodate intense traffic without a dramatic impact on travel times: during the traffic peak hours, drivers tend to choose alternate routes and exploit minor roads, spreading the traffic over the whole street layout. Thus drivers avoid increasing their travel times – and cell residence times – by better leveraging the road capacity during rush hours. In absolute terms, the CDFs highlight the short residence time one can expect in vehicular environments: 70% of the users spend less than one minute in a cell, and half of them leave the cell in 30 seconds at most. Less than 5% of the vehicles spend more than two minutes in the same cell. These numbers are clearly due to the elevate mean speed.

As it was the case for inter-arrival times, we investigate whether a same law can describe the empirical cell residence time distributions in the case of pervasive vehicular access. Fig. 2.12(b) and Fig. 2.12(c) display representative samples of the theoretical distribution fittings, again at 8:00 pm. We found the Weibull distribution to provide the best fit. However, the Weibull distribution fitting always showed an exponent close to one, i.e., it tended to an exponential law. It is thus unsurprising that the exponential curve also fits quite well the empirical data in the plots.

These results highlight how *the common exponential assumption may hold for cell residence times generated by vehicular access*. Moreover, a single distribution instance may be sufficient to describe the overall cell residence time behavior, since *cell residence times tends not to vary significantly over the whole day*. Finally, we stress that *vehicular access generates remarkably low residence times*. This latter observation, together with our previous analysis of the high pervasive vehicular access load in Sec. 2.3.1, indicates that guaranteeing high capacity while limiting the handover cost risks to be a complex task in pervasive vehicular access.

2.4 Conclusions of the Chapter

In this Chapter, I presented results relevant to the characterization of fundamental properties of vehicular networks. Overall, these results yield several contributions, among which one can highlight: (i) the generation of a large-scale urban vehicular mobility dataset that represents today's state of the art and that has been widely employed by the research community during the past couple of years⁹; (ii) the demonstration that the level of realism of other vehicular mobility datasets commonly employed in the literature is not sufficient for networking research; (iii) the characterization of the poor availability, reliability and navigability of vehicular networks based on a pure ad hoc paradigm, hence the identification of the limits of vehicle-to-vehicle communication in terms of services it can support; (iv) the definition of several major implications that a pervasive cellular access from moving vehicles would have on the radio access network.

Bibliography

- [1] F. Bai, N. Sadagopan, A. Helmy, "The IMPORTANT framework for analyzing the Impact of Mobility on Performance Of Routing protocols for Adhoc NeTworks", *Elsevier Ad Hoc Networks*, vol.1, 2003.
- [2] M. Fiore, J. Härri, "The Networking Shape of Vehicular Mobility", *ACM MobiHoc*, Hong Kong, PRC, May 2008.
- [3] W. Viriyasitavat, F. Bai, O.K. Tonguz, "Dynamics of Network Connectivity in Urban Vehicular Networks", *IEEE Journal on Selected Areas in Communications*, vol.29, no.3, pp.515–533, 2011.
- [4] A. K. Saha, D. B. Johnson, "Modeling Mobility for Vehicular Ad Hoc Networks", *ACM VANET*, Philadelphia, PA, USA, Oct. 2004.
- [5] S. Jaap, M. Bechler, L. Wolf, "Evaluation of Routing Protocols for Vehicular Ad Hoc Networks in City Traffic Scenarios", *IEEE ITSC*, Vienna, Austria, Sep. 2005.
- [6] D. Krajzewicz, G. Hertkorn, C. Rossel, P. Wagner, "SUMO (Simulation of Urban MObility): An open-source traffic simulation", *SCS MESM*, Sharjah, United Arab Emirates, Sep. 2002.
- [7] J. Härri, M. Fiore, F. Filali, C. Bonnet, "Vehicular mobility simulation with VanetMobiSim", *Simulation*, vol.87, no.4, Apr. 2011.
- [8] J. Härri, F. Filali, C. Bonnet, "Mobility models for vehicular ad hoc networks: a survey and taxonomy", *IEEE Communications Surveys and Tutorials*, vol.11, no.4, pp.19–41, Dec. 2009.

⁹A non-comprehensive list of universities and research centers using our dataset includes: KIT, Germany; EURECOM, France; IMDEA Networks, Spain; Universidade Federal de Minas Gerais, Brazil; Ilmenau University of Technology, Germany; University of Luxembourg, Luxembourg; Tampere University of Technology, Finland; University of Cyprus, Cyprus.

- [9] H. Conceição, L. Damas, M. Ferreira, J. Barros, “Large-Scale Simulation of V2V Environments,” *ACM SAC*, Fortaleza, Brazil, Mar. 2008.
- [10] C. Barberis, G. Malnati, “Epidemic information diffusion in realistic vehicular network mobility scenarios,” *IEEE ICUMT*, St. Petersburg, Russia, Oct. 2009.
- [11] M. Fiore, C. Casetti, C.-F. Chiasserini, P. Papadimitratos, “Discovery and Verification of Neighbor Positions in Mobile Ad Hoc Networks,” *IEEE Transactions on Mobile Computing*, to appear.
- [12] R. Baumann, F. Legendre, P. Sommer, “Generic mobility simulation framework (GMSF),” *ACM MobilityModels*, Hong Kong, PRC, Apr. 2008.
- [13] C. Lochert, H. Hartenstein, J. Tian, H. Fuessler, D. Hermann, M. Mauve, “A routing strategy for vehicular ad hoc networks in city environments,” *IEEE IV*, Columbus, OH, USA, Jun. 2003.
- [14] M. Ferreira, H. Conceição, R. Fernandes, O.K. Tonguz, “Stereoscopic aerial photography: an alternative to model-based urban mobility approaches,” *ACM VANET*, Beijing, PRC, Sept. 2009.
- [15] G.S. Thakur, P. Hui, A. Helmy “Modeling and Characterization of Urban Vehicular Mobility using Web Cameras,” *NetSciCom*, Orlando, FL, USA, Mar, 2012.
- [16] F. Bai, B. Krishnamachari, “Spatio-temporal variations of vehicle traffic in VANETs: facts and implications,” *ACM VANET*, Beijing, PRC, Sep. 2009.
- [17] iTetris, <http://ict-itetris.eu>.
- [18] Y. Pigne, G. Danoy, P. Bouvry, “A Vehicular Mobility Model based on Real Traffic Counting Data,” *Nets4Cars*, Oberpfaffenhofen, Germany, Mar. 2011.
- [19] P. Wagner, K. Nagel, “Microscopic modeling of travel-deman: Approaching the home-to-work problem,” *Transportation Research Board Annual Meeting*, Washington, DC, USA, Jan. 1999.
- [20] M. Caliskan, B. Rech, A. Lübke, M. Mauve, “Collection of Dedicated Information in Vehicular Ad Hoc Networks,” *World Congress on Intelligent Transport Systems*, S. Francisco, CA, USA, Nov. 2005.
- [21] B. Raney, N. Cetin, A. Völlmy, M. Vrtic, K. Axhausen, K. Nagel, “An agent-based microsimulation model of Swiss travel: First results,” *Networks and Spatial Economics*, vol.3, no.1, pp.23–41, Jan. 2003.
- [22] N. Cetin, A. Burri, K. Nagel, “A large-scale multi-agent traffic microsimulation based on queue model,” *STRC*, Ascona, Switzerland, Mar. 2003.
- [23] J.G. Jetcheva, Y.-C. Hu, S. PalChaudhuri, A.K. Saha, D.B. Johnson, “Design and evaluation of a metropolitan area multitier wireless ad hoc network architecture,” *IEEE WMCSA*, Monterey, CA, USA, Oct. 2003.

- [24] M. Doering, T. Pögel, W.-B. Pöttner, L. Wolf, “A new mobility trace for realistic large-scale simulation of bus-based DTNs,” *ACM CHANTS*, Chicago, IL, USA, Sep. 2010.
- [25] J. Burgess, B. Gallagher, D. Jensen, B.N. Levine, “MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks,” *IEEE INFOCOM*, Barcelona, Spain, Apr. 2006.
- [26] H.-Y. Huang, P.-E. Luo, M. Li, D. Li, X. Li, W. Shu, M.-Y. Wu, “Performance Evaluation of SUVnet With Real-Time Traffic Data,” *IEEE Transactions on Vehicular Technology*, vol.56, no.6, pp.3381–3396, Nov. 2007.
- [27] M. Sede, X. Li, D. Li, M.-Y. Wu, M. Li, W. Shu, “Routing in Large-Scale Buses Ad Hoc Networks,” *IEEE WCNC*, Las Vegas, NV, USA, Mar. 2008.
- [28] L. Zhang, J. Ahn, B. Krishnamachari, F. Bai, “Optimizing Content Dissemination in Heterogeneous Vehicular Networks,” *PATH/Tsinghua Workshop on ITS*, Apr. 2010.
- [29] J. Yuan, Y. Zheng, X. Xie, G. Sun, “Driving with knowledge from the physical world,” *ACM SIGKDD*, San Diego, CA, USA, Aug. 2011.
- [30] Cab Spotting, <http://cabspotting.org/api>.
- [31] sim^{TD}, <http://www.simtd.de>.
- [32] OpenStreetMap, <http://www.openstreetmap.org>.
- [33] Simulation of Urban Mobility, <http://sumo.sourceforge.net>.
- [34] S. Krauß, P. Wagner, C. Gawron, “Metastable States in a Microscopic Model of Traffic Flow”, *Physical Review E*, vol.55, no.304, pp.55–97, May 1997.
- [35] D. Krajzewicz, “Kombination von taktischen und strategischen Einflüssen in einer mikroskopischen Verkehrsflusssimulation”, in T. Jürgensohn, H. Kolrep (editors), *Fahrermodellierung in Wissenschaft und Wirtschaft*, VDI-Verlag, pp.104–115, Berlin-Adlershof, Germany, 2009.
- [36] G. Rindsfüser, J. Ansorge, H. Mühlhans, “Aktivitätenvorhaben”, in K.J. Beckmann (editor), *SimVV Mobilität verstehen und lenken – zu einer integrierten quantitativen Gesamtsicht und Mikrosimulation von Verkehr*, Final report, Ministry of School, Science and Research of Nordrhein-Westfalen, Düsseldorf, Germany, 2002.
- [37] M. Ehling, W. Bihler, “Zeit im Blickfeld. Ergebnisse einer repräsentativen Zeitbudgeterhebung”, in K. Blanke, M. Ehling, N. Schwarz (editors), *Schriftenreihe des Bundesministeriums für Familie, Senioren, Frauen und Jugend*, vol.121, pp.237–274, Kohlhammer, Stuttgart, Germany, 1996.

- [38] C. Varschen, P. Wagner, “Mikroskopische Modellierung der Personenverkehrsnachfrage auf Basis von Zeitverwendungstagebüchern”, *Stadt Region Land*, vol.81, pp.63–69, 2006.
- [39] C. Gawron, “An Iterative Algorithm to Determine the Dynamic User Equilibrium in a Traffic Simulation Model”, *International Journal of Modern Physics C*, vol.9, no.3, pp.393–407, 1998.
- [40] Bezirksregierung Köln, “Luftreinhalteplan für das Stadtgebiet Köln – Erste Fortschreibung 2012”, 2012.
- [41] M. Kafsi, P. Papadimitratos, O. Dousse, T. Alpcan, J.-P. Hubaux, “VANET Connectivity Analysis”, *IEEE AutoNet*, New Orleans, LA, USA, Dec. 2008.
- [42] G. Pallis, D. Katsaros, M. D. Dikaiakos, N. Loulloudes, L. Tassiulas, “On the Structure and Evolution of Vehicular Networks”, *IEEE/ACM MASCOTS*, London, UK, Sep. 2009.
- [43] W. Viriyasitavat, F. Bai, O.K. Tonguz, “Dynamics of Network Connectivity in Urban Vehicular Networks”, *IEEE Journal on Selected Areas in Communications*, 29(3):515–533, Mar. 2011.
- [44] R. Albert, A. Barabási, “Statistical Mechanics of Complex Networks”, *Reviews of Modern Physics*, 74(1):47–97, Jan. 2002.
- [45] M.E.J. Newman, “The Structure and Function of Complex Networks”, *SIAM Review*, 45(2):167–256, 2003.
- [46] S Boccaletti, V. Latora, Y. Moreno, M. Chavez, D. Hwang, “Complex Networks: Structure and Dynamics”, *Physics Reports*, 424(4-5):175–308, 2006.
- [47] P. Crucitti, V. Latora, S. Porta, “Centrality measures in spatial networks of urban streets”, *Physics Review E*, 73(3), Mar. 2006.
- [48] Cisco, “Global Mobile Data Traffic Forecast Update, 20122017”, White Paper, Feb. 2013.
- [49] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, S. Madden, “A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks”, *ACM MobiCom*, Las Vegas, NV, USA, Sept. 2006.
- [50] Mobilfunk in Köln, <http://www.fst-gsm.de/>.
- [51] S. Karpinski, E.M. Belding-Royer, K.C. Almeroth, “Living on the Edge: The Distribution of Flows Across Mobile Nodes in Large Wireless Networks,” *Technical Report*, 2005.

Chapter 3

Data transfers in vehicular networks

The transfer of data to and from vehicles is today performed by networked vehicles through traditional cellular communication. Although simple to implement, this approach is expected not to scale well in the future, due to the growth in the demand of vehicular users and the difficulties of the cellular paradigm to keep the pace with the overall increase of mobile data traffic. There is thus a need to identify networking paradigms that can complement cellular access and offload the existing infrastructure. Distributed techniques that exploit local, direct communication among vehicles represent a possible solution to that problem.

This Chapter includes results on the management of data transfers between vehicles and the radio access infrastructure that take advantage of vehicle-to-vehicle communication. In Sec. 3.1, I will discuss the downlink direction, addressing the problem of download of large contents stored on Internet-based servers by moving vehicles. The main objective there will be understanding how to use vehicle-to-vehicle communication to improve the system performance and reduce the load on the network infrastructure. In Sec. 3.2, I will consider instead the uplink case, i.e., the upload of small-sized but frequently-updated data from vehicles to the cellular infrastructure. Again, the goal will be to explore the potential advantage provided by the use of vehicle-to-vehicle communication as a complement to the traditional cellular communication.

The results presented in the following are the fruit of joint works with Francesco Malandrino, Razvan Stanica, Oscar Trullols-Cruces, Carla-Fabiana Chiasserini, Claudio Casetti, and Josè Maria Barcelo-Ordinas. The related publications are as follow:

- M. Fiore, J.M. Barcelo-Ordinas, “*Cooperative download in urban vehicular networks*,” IEEE MASS, Macau, China, October 2009.
- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, “*Content Downloading in Vehicular Networks: What Really Matters*,” IEEE INFOCOM Miniconference, Shanghai, China, April 2011.
- O. Trullols-Cruces, M. Fiore, J.M. Barcelo-Ordinas, “*Cooperative download in vehicular environments*,” IEEE Transactions on Mobile Computing, 11(4), April 2012.

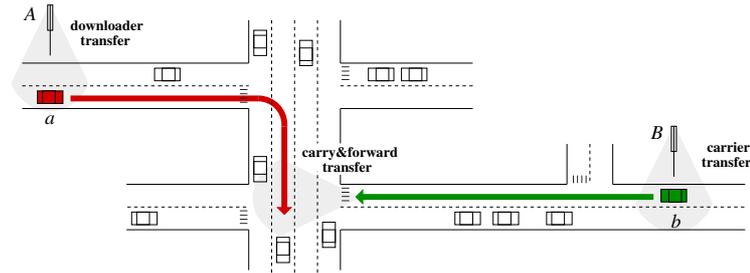


Figure 3.1: Cooperative download scenario involving DSRC-enabled vehicles.

- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, *Offloading Cellular Networks through ITS Content Download*, IEEE SECON, Seoul, Korea, June 2012.
- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, *“Optimal Content Downloading in Vehicular Networks,”* IEEE Transactions on Mobile Computing, 12(7), July 2013.
- R. Stanica, M. Fiore, F. Malandrino, *“Offloading Floating Car Data,”* IEEE WoWMoM, Madrid, Spain, June 2013.
- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, *“Content Download in Vehicular Networks in Presence of Noisy Mobility Prediction,”* IEEE Transactions on Mobile Computing, to appear.

3.1 The downlink direction: cooperative download

As previously pointed out in Sec. 2.3, the emergence of networked vehicles will lead users moving on cars and public transportations to expect that every information content is always readily available onto their tablets, smartphones, or onto the integrated on-board communication and computing systems of their vehicle. However, the success of mobile services has recently led to a surge in mobile data traffic: typical 3G users have been shown to generate between 50 and 120 times the traffic of a basic-feature phone, and the cellular access infrastructure has been mostly used for data rather than voice for several years now [1]. According to Cisco projections, the global mobile data traffic will still grow at a compound annual rate of 66% in the next four years, leading to an overall 13-fold increase between 2012 and 2017 [2]. The same study underlines that the access network connection speed will not keep the pace, augmenting just 7-fold by the same date. Moreover, the bandwidth availability in the cellular backhaul will also be hard-pressed to keep the pace with what is being offered on the air interface [3].

As a result, many observers call for the development of alternative communication systems to support and relieve the congested cellular network in areas where the demand by mobile users is expected to be the thickest. In particular, large-sized content downloading accounts for most of the traffic in access networks [4], and is thus a prime candidate for offloading.

In order to design a network architecture that will scale to support the mass of vehicular users, one possibility is to offload part of the traffic to Dedicated Short-Range Communication (DSRC), through direct infrastructure-to-vehicle (I2V) transfer, as well as vehicle-to-vehicle (V2V) data relaying. Specifically, we consider a scenario where users aboard cars can exploit an infrastructure of roadside units (RSUs) to access the Internet-based servers that host the desired contents. The coverage provided by such an infrastructure is intermittent, since a pervasive deployment of RSUs dedicated to vehicular access is impractical for economic and technical reasons. We can however compensate for holes in the RSU coverage by complementing I2V communication with opportunistic V2V transfers. An example of such a scenario is shown in Fig. 3.1, where vehicle a downloads part of some content from RSU A , while RSU B delegates another portion of the same content to the passing-by vehicle b . When b encounters a , V2V communication is leveraged to transfer to a the data carried by b . We refer to such an approach as cooperative downloading¹.

3.1.1 Drawing the performance bounds

We first try to quantify the actual potential of cooperative downloading, by identifying the performance limits achievable through DSRC-based I2V/V2V communication. To this end, we assume ideal conditions from a system engineering viewpoint, i.e., the availability of preemptive knowledge of vehicular trajectories and perfect scheduling of data transmissions. We also assume that we can plan the RSU infrastructure deployment and enforce the data transfer paradigm to be employed by each vehicle for every content chunk. By leveraging such degrees of freedom, we cast the downloading process to a mixed integer linear programming (MILP) max-flow problem.

Problem formulation

The objective of the MILP optimization problem is to maximize the aggregate throughput², attained by the cooperative downloading system. To this aim, we have to jointly solve two problems: (i) given a set of candidate locations and a number of RSUs to be activated, we need to identify the deployment yielding the maximum throughput; (ii) given the availability of different data transfer paradigms, possibly involving relays, we have to determine how to use them in order to maximize the data flow from the infrastructure to the downloaders.

Our approach starts from a road layout with a set of candidate RSU deployment locations, and an associated vehicular mobility dataset. From those, we build a time-expanded weighted graph [6], hereinafter dynamic network topology graph

¹We remark that the cooperative downloading scenario is consistent with the assumption that on-board users will behave similarly to home users: not all on-board users will download large files all the time, exactly as it happens in today's wired networks, where the portion of queries for large contents is small [5].

²Any attempt at including fairness in the formulation (e.g., via *max-min* approaches) failed due to the heterogeneity that characterizes road traffic over typical street layouts. Indeed, providing significant throughput to vehicles traveling along minor roads proved impossible via DSRC only.

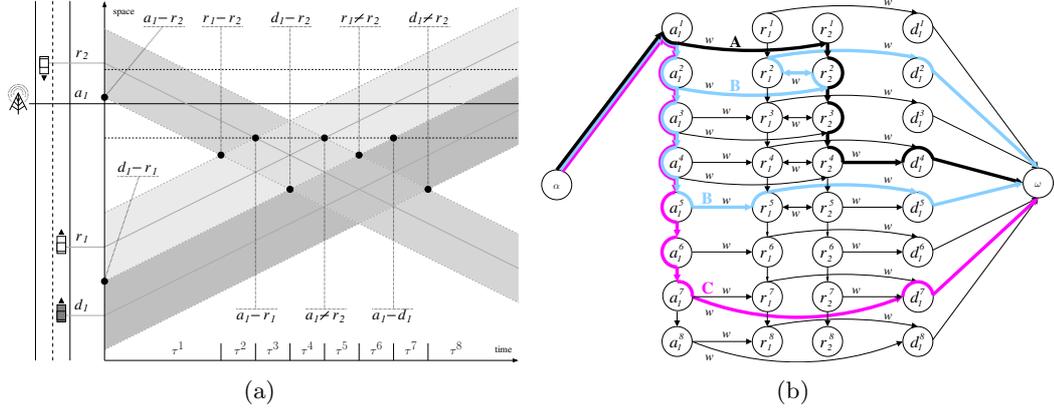


Figure 3.2: Dynamic network topology graph (DNTG) sample. (a) Contact events among vehicles and one candidate RSU location, with shadowed areas representing halved transmission ranges, so that links exist when two shadowed areas touch or overlap, and break when such areas become disjoint. (b) Corresponding DNTG.

(DNTG), that represents the temporal evolution of the vehicular network. An intuitive example of DNTG derivation is given in Fig. 3.2, in presence of one RSU candidate location and three vehicles d_1, r_1, r_2 , with d_1 being a downloader and r_1, r_2 possibly acting as relays. Fig. 3.2(a) depicts the spatio-temporal evolution of node positions: there, contact events are highlighted through the times at which links are established or lost. The durations of the time frames within which the network connectivity is unchanged³, are denoted by τ^1, \dots, τ^8 . In Fig. 3.2(b), frames correspond to rows of vertices in the DNTG, where intra-nodal edges connect vertices representing the same vehicle or candidate RSU location over time. Note that the graph is completed by the virtual vertices α and ω , needed to formulate the max-flow problem. The DNTG models all possible opportunities through which data can flow from the RSUs to the downloaders, via the different data transfer paradigms enabled by DSRC-based I2V/V2V communication: (i) *direct download* from the RSU to the downloader, as in path C; (ii) *connected forwarding* from the RSU to the downloader through a multi-hop path of relays, as in paths B (of 3 hops, in frame 2, and 2 hops, in frame 5, respectively); (iii) *carry-and-forward* by one or multiple intermediate relays that store and physically carry the data before the latter is delivered to the downloader, as in path A.

The DNTG is used to formulate a max-flow problem, whose solution yields both the optimal RSU deployment over a given road layout and the optimal combination of any possible I2V and V2V data transfer paradigm. It thus represents the theoretical upper bound to the downloading throughput in our scenario. Namely, we maximize the flow of data⁴ from α to ω , i.e., the total amount of downloaded data. Over a

³In this example we neglect the impact of finer aspects, such as, e.g., the achievable network-layer rate, on the network connectivity. These are however accounted for in the performance evaluation.

⁴We do not model the scheduling of the single packets transmitted within each frame, but consider data as a continuous flow. Such an approximation has a negligible impact on large-scale transfers.

time window of F frames, our objective can be expressed as

$$\max \sum_{k=1}^F \sum_{v_i^k \in \mathcal{D}^k} x(v_i^k, \omega), \quad (3.1)$$

where v_i^k is a vehicle from the set \mathcal{D}^k of downloaders at frame k , and $x(\cdot, \cdot)$ denotes the traffic flow over an edge connecting two generic vertices.

The max-flow problem needs to be solved taking into account several constraints, listed below⁵.

- Non-negative flow. The flow on every existing DNTG edge must be greater than or equal to zero.
- Flow conservation. For any vertex in the DNTG, the amount of incoming flow must equal the amount of outgoing flow.
- Channel access. We assume that vehicles and RSUs are equipped with a radio interface using an IEEE 802.11-based MAC protocol. Different frequencies are reused so that adjacent RSUs operate on separate channels, and one channel is dedicated to V2V communication only. Also, for speed and reliability all transmissions are unicast. These settings imply that, given a vehicle, none of the following events can take place simultaneously, and the time span of each frame must be shared among them:
 1. the vehicle transmits to a neighboring vehicle;
 2. a neighboring vehicle receives from any relay;
 3. the vehicle receives from a neighboring relay;
 4. a neighboring relay transmits to any vehicle;
 5. the vehicle receives from a neighboring RSU.
- Overlapping RSU coverage. Since RSUs operate on different frequency channels, a vehicle within coverage of two or more RSUs can only communicate with one of them during the same time frame.
- Maximum number of RSUs. We impose that a maximum number A of candidate locations are selected for RSU deployment.

Implicitly, our constraints also enforce that vehicular users are rational, i.e., they can be engaged in relaying traffic for others only if they are not currently retrieving some content for themselves.

Unfortunately, the resulting MILP problem is too complex to be solved for large instances (e.g., large geographical areas, high number of vehicles participating in the content downloading, or large number of candidate RSU locations). To be able to analyze such cases, we resort to a graph sampling approach, by: (1) sampling the DNTG obtaining a small, yet representative, sub-graph, which includes all relevant candidate RSU locations; (2) finding the optimal RSU deployment using such a sub-graph; (3) applying the obtained deployment to the full graph and optimizing the flows – a linear programming (LP) problem that can be easily solved as it does not involve Boolean variables.

⁵For the sake of readability, I decided not to detail the rather complex mathematical formulation of the constraints and the associated notation. The complete formulation can be found in the paper appeared on IEEE Transactions on Mobile Computing 12(7).

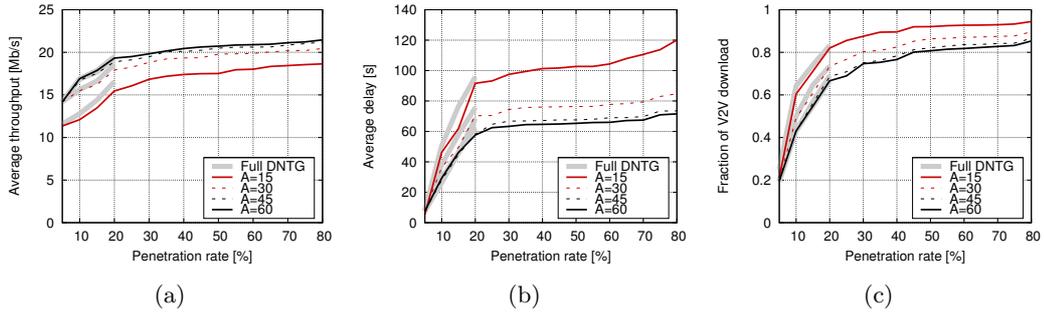


Figure 3.3: Cooperative download: average per-downloader metrics obtained from the max-flow solution, imposing a 2-hop limit to relaying. Results are shown as a function of the technology penetration rate. (a) Throughput. (b) Delay. (c) V2V downloading fraction.

Optimal performance

We evaluate the optimal performance of the cooperative downloading system in a typical urban vehicular environment, modeling network-layer rates according to the experimental results in [7]. Our evaluation is run on the basis of several varying parameters: (1) the technology penetration rate p , i.e., the fraction of vehicles equipped with a radio interface and participating in the downloading process; (2) the percentage of such communication-enabled vehicles that concurrently request content, i.e., that act as downloaders; (3) the number of deployed RSUs. Unless otherwise specified, we will consider $d = 0.01$ (i.e., 1% of the vehicles participating in the network) – a reasonable value as observed in wired networks [5].

Fig. 3.3 portrays the evolution of key performance indicators when the technology penetration rate p (in abscissa) varies between 5% and 80%. The thin curves refer to the results obtained with the sampling approach for different extensions of the roadside infrastructure, ranging from $A = 15$ to $A = 60$ RSUs (mapping to densities from 0.75 to 3 RSUs per km^2), with the latter value achieving a complete coverage of the road topology. Thicker, light-grey lines represent instead the results derived by solving the full DNTG problem, whenever possible, and show good agreement with those drawn through sampling. Here, we also limit the relaying to 2 hops (i.e., we allow at most one relay to intervene between the RSU and the downloader), an assumption that we will relax later on.

The average per-downloader throughput, in Fig. 3.3(a), is good in all conditions, scoring well above 10 Mb/s even in low- p , low- A scenarios, and more than 20 Mb/s in presence of a wide I2V and V2V technology adoption. Delays, in Fig. 3.3(b), are in the order of tens of seconds, a good result when considering the delay-tolerant nature of most V2V transfers. Fig. 3.3(c) highlights the fundamental role of V2V traffic relaying in attaining such positive performance: most of the content is indeed received by downloaders from relays rather than directly from RSUs.

When separating the effects of A and p for each metric, we remark that the availability of a more pervasive infrastructure helps, although less than one could expect. Indeed, a pervasive 60-RSU deployment only results in a 3-Mb/s throughput

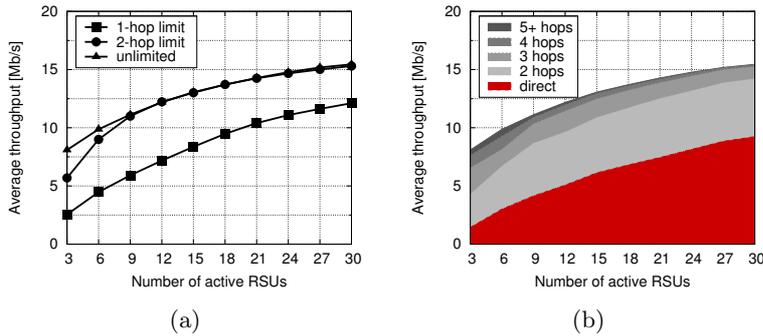


Figure 3.4: Cooperative download at low-penetration regime. (a) Average per-downloader throughput. (b) Number of transfer hops split-up of the average per-downloader throughput, when the number of relays is unconstrained.

gain over a simple 15-RSU deployment, and a negligible 10-s delay gain over a 30-RSU deployment. Higher improvements can be instead obtained from the spread of in-car communication interfaces, with an average throughput increase of 8 Mb/s as p grows from 5% to 80%. More importantly, all performance metrics are consistent in revealing the existence of two operation regimes that depend on p . The first, when $p < 20\%$, i.e., at early stages of the technology adoption, is characterized by a very strong dependency on direct I2V communication, which implies lower throughput, lower delays, and unfairness between vehicles traveling on high-traffic arteries and minor roads⁶ (with the former experiencing much better downloading performance than the latter). The second, for $p > 30\%$, i.e., in presence of a quite mature technology, features instead a jump in throughput, delay and fairness, due to a dramatic increase in the use of V2V relaying. As the impact of the system settings is different within these low- and high-penetration regimes, we will study them separately⁷.

We first employ the low- p regime to relax the assumption on the 2-hop limit in relaying, and compare three different scenarios, where we allow (i) only direct I2V communication, (ii) a 2-hop limited relaying, and (iii) unlimited relaying. Fig. 3.4(a) depicts the average throughput achieved in the three cases, when $p = 10\%$. In the direct case, downloaders rely on direct contacts with the RSUs only, which leads to a significantly lower throughput than in the other cases. In fact, allowing a single relay between RSUs and downloaders yields a throughput gain between 20% and 100%, depending on the infrastructure coverage. Even more interestingly, considering transfers over 3 hops or more yields almost no advantage over the case where a 2-hop limit is enforced. In order to explain the latter effect, we fragment the downloading throughput measured in the unlimited relaying scenario, according to the number of hops traveled by packets to reach their destination. Fig. 3.4(b) shows how, even when unlimited hops are allowed, a large majority of the relayed data

⁶Results on the fairness are not shown here for the sake of brevity. They can be retrieved in the paper appeared on IEEE Transactions on Mobile Computing 12(7).

⁷Due to the different problem size, we will employ the max-flow problem solution on the complete and on the sampled graph in the low- and high-penetration regime, respectively.

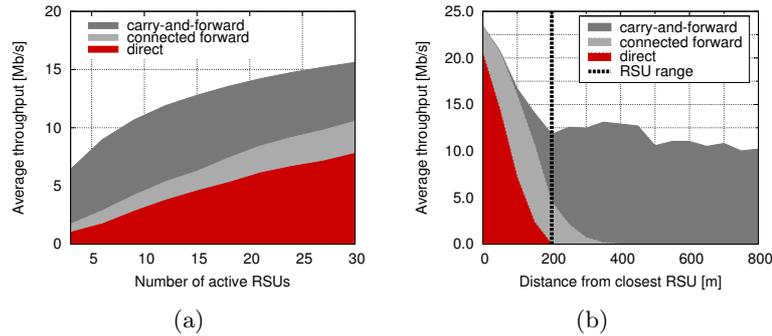


Figure 3.5: Cooperative download at low-penetration regime. (a) Transfer paradigm split-up of the average per-downloader throughput. (b) Average per-downloader throughput versus the distance between the downloader and the closest RSU.

traffic arrives at destination in just two hops: it is thus clear that the availability of additional hops, while granting more flexibility to the max-flow problem solution, only leads to minor adjustments that have a negligible impact on the overall downloading performance.

We thus revert to the 2-hop relaying case, and explore the performance from a different viewpoint, i.e., that of the contribution of the different transfer paradigms to the overall throughput, in Fig. 3.5(a). There, we can observe that the fraction of content downloaded through V2V relaying is predominantly due to carry-and-forward and somewhat constant, contributing approximately 5 Mb/s to the overall throughput. The fact that the V2V contribution is not affected by the number of RSUs deployed is counterintuitive, but can be understood by observing the average per-downloader throughput as a function of the downloader distance from the closest RSU, in Fig. 3.5(b). Clearly, direct I2V transfers can only occur within the RSU transmission range (200 m according to [21], and highlighted by the vertical line on the plot), while 2-hop connected forwarding reaches at most twice such a distance. Carry-and-forward is not distance bounded, hence it can reach downloaders that are very far from RSUs. The key observation is that V2V relaying frequently occurs within range of RSUs. In fact, at a distance of 100 m, i.e., half of the maximum transmission range of the RSU, communication largely takes place through relays. The reason is that our model realistically accounts for the fact that the network-layer rate decreases with distance, which can make the use of high-rate multi-hop paths preferable to low-rate direct transfers. This explains why, even in presence of a pervasive RSU coverage, relaying is employed to improve the wireless resource utilization, and thus the overall throughput.

The high- p regime maps to the widespread adoption of I2V and V2V communications, a condition that is likely to be associated with a growth in the downloading activity of the average vehicular user. Evaluating the impact of the amount of concurrent downloaders, i.e., users requesting some content during the same time interval, becomes then critical. Fig. 3.6 shows the impact of the concurrent downloader fraction d on the performance metrics, when $p = 50\%$. We consider d ranging

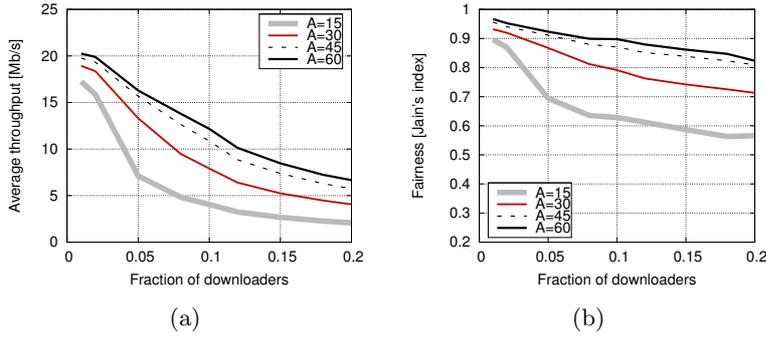


Figure 3.6: Cooperative download at high-penetration regime, versus the fraction of concurrent downloaders. (a) Average per-downloader throughput. (b) Fairness.

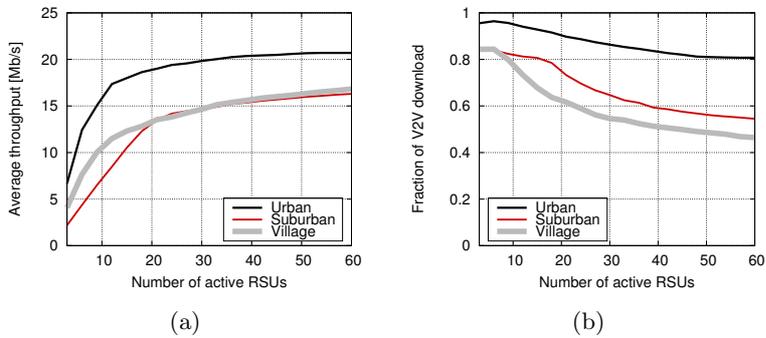


Figure 3.7: Cooperative download at high-penetration regime, in diversely urbanized areas. (a) Average per-downloader throughput. (b) V2V downloading fraction.

between 0.01 and 0.2, the latter representing a highly-loaded system, in which one out of five users is downloading data at any given time instant. As one could expect, when the system load grows, increasing the number of RSUs comes in handy, and can noticeably improve throughput, in Fig. 3.6(a), and fairness, in Fig. 3.6(b). Also, increasing the demand (especially, for $d > 0.1$) reduces the per-user throughput, due to the augmented contention for the limited wireless resources. Less intuitively, fairness degrades as d grows: when the number of simultaneous downloaders increases, vehicles on secondary roads experience less channel contention, hence higher throughput than vehicles traveling on main (and more crowded) roads.

The pervasive presence of communication-enabled vehicles is also expected to extend V2V-based services to non-urban areas. We are thus interested in understanding the impact of different urbanization levels on the downloading performance at the high- p regime. Fig. 3.7 shows the throughput and V2V download fraction in presence of representative urban, suburban and rural scenarios, featuring diverse road layouts and traffic levels. The road topology has a major impact on the downloading process, and the highest throughput, in Fig. 3.7(a), is achieved in the urban scenario, while drivers in the suburban and rural environments experience similarly worse performance. In the urban scenario, many vehicles travel over different roads,

which basically allows a spatial reuse of the wireless medium. The rural scenario is similar to the urban one in that vehicular traffic is quite evenly spread over the road topology; however, the lower number of vehicles reduces the availability of relays, as also evident from Fig. 3.7(b). In the suburban scenario, instead, a high vehicular density is concentrated on a few roads: the consequent channel congestion yields reduced per-downloader throughput.

3.1.2 Practical solutions

Modeling the downloading process as a max-flow optimization problem allowed us to outline the performance limits of a vehicular content downloading system. The promising results we obtained pushed us one step forward, relaxing the assumption of complete knowledge of the system and designing practical mechanisms capable of approaching the performance bounds.

To that end, we face a scheduling problem at RSUs, i.e., deciding at each time instant which content chunk should be transferred to which passing-by (downloader or relay) vehicle. Next, we propose two possible solutions that let each RSU either (i) solve an optimization problem based on a probabilistic graph model of the cooperative download system, or (ii) run probabilistic heuristics leveraging so-called contact maps.

Probabilistic graph-based scheduling

We assume that each RSU can access a prediction of the future vehicular mobility in the region, e.g., from an Internet-based traffic control center. We do not make any assumption on the precise nature of the prediction technique, but represent it in terms of its accuracy in anticipating contacts between vehicles and RSUs, and among vehicles themselves⁸. The prediction is used by each RSU to build a time-expanded weighted graph, in a very similar to what done for the DNTG in Sec 3.1.1: the only difference is that weights on edges are now associated with a probability, since contacts are issued by a prediction and are not certain anymore.

The graph is then exploited to formulate a non-integer linear programming (LP) optimization problem, whose aim is to maximize the amount of data reaching the downloaders. In order to make the solution more realistic, we impose a maximum time limit T to the cooperative download of a content: once expired, the downloader will revert to costly cellular transfers to retrieve the remaining portion of the content. Thus, each RSU formulates the optimization problem as

$$\max \sum_m \sum_c \sum_{k=t_{m,c}}^{t_{m,c}+T} \phi_{m,c}^k. \quad (3.2)$$

There, three sums are over all downloaders v_m 's, all content c 's that v_m has requested, and the time frame k 's, respectively. By $t_{m,c}$, we denote the time which downloader

⁸We adopt a so-called fog-of-war model, which capture different levels of prediction accuracy and whose details are omitted here. For a thorough discussion of the prediction accuracy model, including its validation against actual Markovian models, please refer to the paper to appear on IEEE Transactions on Mobile Computing.

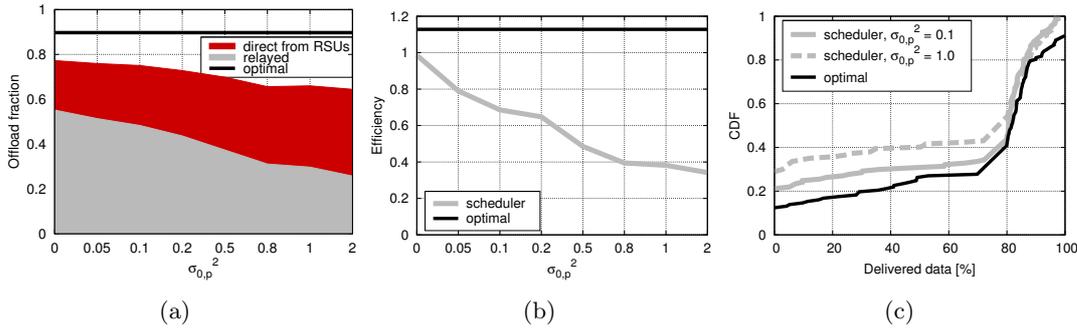


Figure 3.8: Cooperative download via a probabilistic graph-based scheduling, versus the vehicular mobility prediction accuracy. (a) Fraction of demanded content that is retrieved through cooperative download. (b) Efficiency of the cooperative download.

v_m generated a request for content c , while $\phi_{m,c}^k$ is the fraction of the content that m downloads at frame k through the vehicular network, computed as

$$\phi_{m,c}^k = \frac{1}{S_c} \left[\sum_{(v_i^k, v_m^k) \in \mathcal{E}_v^k} f_c(v_i^k, v_m^k) + \sum_{(r_i^k, v_m^k) \in \mathcal{E}_r^k} f_c(r_i^k, v_m^k) \right]. \quad (3.3)$$

In (3.3), S_c is the content size while $f_c(\cdot, \cdot)$ is the expected flow for content c over the edge between two vertices. The first sum is over all vehicle-to-vehicle edges at frame k , forming the set \mathcal{E}_v^k . The second sum is over all RSU-to-vehicle edges at frame k , forming the set \mathcal{E}_r^k . Thus, $\phi_{m,c}^k$ represents the content flow at frame k over the I2V and V2V edges, and thus v_i and r_i denote, respectively, a relay and an RSU storing at frame k part of, or all, content c requested by v_m . As in the case of the optimal solution discussed in Sec 3.1.1, the problem is completed by several constraints on non-negative flows, flow causality and channel access. An important difference is instead that this formulation falls in the non-integer LP category that allows for polynomial-time solution; in fact, our formulation is even suitable to be solved by each RSU in real time [8].

Fig. 3.8 outlines the performance of the probabilistic graph-based scheduling, in presence of a typical urban scenario. We consider that RSUs are deployed at the most crowded road intersections, with a density of 1 RSU/km². Around 5% of the users concurrently request 10-MB contents with a patience time T of 2 minutes.

In Fig. 3.8(a), we show the offload fraction, i.e., the portion of demanded content that is retrieved through cooperative download (and thus offloaded from the cellular infrastructure). We observe that cooperative download can relieve the cellular infrastructure of 70-80% of the cost of content download. A sizable contribution comes from V2V relaying, bearing between 30 and 60% of the content transfer effort, which confirms that opportunistic transfers are highly beneficial in the offload process. These results are coherent with the optimal performance laid out in Sec 3.1.1: in fact, the overall performance is not too far from the optimal one, which would allow a 90% offload in the considered scenario.

The impact of the accuracy of the contact prediction is shown by varying $\sigma_{0,p}^2$,

i.e., the variance of the noise used to model the error on the establishment of contacts, which proved to be an excellent measure of the prediction reliability. Quite surprisingly, very accurate predictions (low values on the x axis) result in a performance that is just slightly better than that scored by almost random contact estimations (high values of $\sigma_{0,p}^2$). Inaccurate predictions lead however to a reduced contribution of V2V with respect to I2V transfers, as the former drops from 60% to less than 30%. The actual cost of an imprecise contact prediction is revealed by Fig. 3.8(b), which shows the offload efficiency, i.e., the ratio of the amount of data delivered to a downloader to that transmitted by the RSUs (to either downloaders or relays). A low efficiency implies a waste of wireless resources at the RSUs, while a high efficiency means that only useful vehicular-based transfers are performed. The efficiency can be higher than 1, since a relay can download some content (or part of it) and then provide it to multiple downloaders interested in the content. The plot clearly shows that, in order to maintain high offload fractions, the less precise is the information on future contacts, the larger is the amount of data the RSUs have to transfer to relays.

Also in the case of practical scheduling solutions, one should not expect fairness in the cooperative download system. Fig. 3.8(c) shows the CDF of the fraction of content that each downloader can retrieve through I2V and V2V communications. Results are shown for quite accurate ($\sigma_{0,p}^2 = 0.1$) and rather imprecise ($\sigma_{0,p}^2 = 1$) predictions, and benchmarked against the optimal case. The distributions clearly identify two classes of downloaders: those that can get a very small percentage (possibly zero) of the data they request, and those (over 50% of the total) that can obtain almost all (80% or more) of the data through cooperative download. These two categories correspond again to users traveling, respectively, on secondary roads and main thoroughfares. The former are seldom under RSU coverage and experience fewer contacts with (relay) vehicles. Interestingly, the latter do not seem to be affected by $\sigma_{0,p}^2$, as curves are very close for high values on the x axis. On the contrary, the percentage of downloaders unable to get any data is sensibly reduced as the contact estimation precision grows. We can thus conclude that an accurate prediction is most useful to offload traffic destined to for hard-to-reach users.

Contact map-based scheduling

The second solution to the cooperative download scheduling problem differs from the first in that it is self-contained, i.e., runs at RSUs only. As a matter of fact, the solution includes the contact prediction phase as well, and thus it does not need any support from, e.g., an external road traffic center. To that end, the contact prediction process is simplified so as to anticipate potential paths for two-hop data transfers only. Such a design choice is not expected to impact significantly the results: as we observed in Sec. 3.1.1, allowing one intermediate relay already allows to attain performance that are close to the optimal.

The contact prediction is performed by each RSU, by training so-called *contact maps*, which are in turn based on the concept of *phases*. We denote as p_{Aa}^k the k -th *production phase* of vehicle a with respect to RSU A , i.e., the k -th of the disjoint time intervals during which vehicle a can steadily download data from A [9].

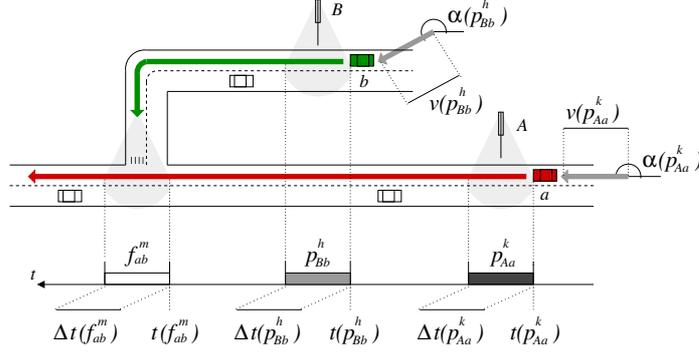


Figure 3.9: Production and forward phase notation.

From a specific RSU perspective, we tag production phases as *local* if they involve that particular RSU: as an example, p_{Bb}^h is a local production phase for RSU B , $\forall b, h$. On the other hand, we label as f_{ab}^m the m -th *forward phase* of vehicle b with respect to vehicle a , i.e., the m -th of the disjoint time intervals during which vehicle b can steadily forward data to vehicle a . Note that production and forward phases do not necessarily correspond to actual data transfers, but just to contacts which could be exploited for data transfers. We also use $t(\cdot)$ to indicate the time at which a production or forward phase starts, and $\Delta t(\cdot)$ to tag its duration. For production phases only, $\alpha(\cdot)$ denotes the global direction of movement of the vehicle at the beginning of the production phase, and $v(\cdot)$ its speed at that same time. The notation is summarized in Fig. 3.9.

A contact map is a data structure that provides an RSU with information on the probability of contact between a vehicle involved in a local production phase and another vehicle. Contact maps can be constructed in a distributed yet consistent way by strictly complying with a set of well defined compatibility rules⁹. With reference to the example in Fig. 3.9, the contacts map at RSU B allows B to know the probability of contact between the local vehicle b and the generic vehicle a . In particular, RSU B knows that b started a local production phase p_{Bb}^h at time $t(p_{Bb}^h)$, while moving with direction $\alpha(p_{Bb}^h)$ and speed $v(p_{Bb}^h)$; also, let us assume B has been informed that a started a production phase p_{Aa}^k with RSU A at time $t(p_{Aa}^k)$, while moving with direction $\alpha(p_{Aa}^k)$ and speed $v(p_{Aa}^k)$. Then, the contacts map at B allows to associate the couple of production phases (p_{Bb}^h, p_{Aa}^k) to historical data on the encounters between vehicles that have previously generated production phases at the two RSUs B and A with timings and mobility similar to those of b and a .

More formally, a contacts map is a set of one-to-one associations between *keys* that encode the significant characteristics of two production phases, and *values*, that store the contacts properties for all couples of production phases that share such characteristics. The key stored at RSU B for two generic production phases p_{Bb}^h and

⁹Details and examples on compatibility rules and contact map construction can be found in the paper appeared on IEEE Transactions on Mobile Computing 11(4).

```

01 set  $\mathbb{P}$  equal to  $\mathbb{P}_{min}$ 
02 for each downloader vehicle  $a$  do
03   if  $a$  is in range of  $B$  do
04     if  $a$  is closer to  $B$  than previous direct downloaders do
05       select  $a$  as destination for direct transfer
06       select no vehicles as carriers for carry&forward transfer
07       set  $\mathbb{P}$  equal to  $\infty$ 
08     done
09   else
10     for each production phase  $p_{Aa}^k$  of  $a$  do
11       until all on-going local production phases are not processed do
12         update delivery potential  $\mathbb{P}_{Aa}^k$ 
13         update carriers list  $\bar{c}_{Aa}^k$ 
14       done
15       if  $\mathbb{P}_{Aa}^k$  is the highest potential computed for  $a$  do
16         set  $\mathbb{P}_a$  equal to  $\mathbb{P}_{Aa}^k$ 
17         set  $\bar{c}_a$  equal to  $\bar{c}_{Aa}^k$ 
18       done
19     done
20   if  $\mathbb{P}_a$  is strictly higher than  $\mathbb{P}$  do
21     select  $a$  as destination for carry&forward transfer
22     select vehicles in  $\bar{c}_a$  as carriers for carry&forward transfer
23     set  $\mathbb{P}$  equal to  $\mathbb{P}_a$ 
24   done
25 done
26 done

```

Figure 3.10: Pseudocode for relay selection at RSU B

p_{Aa}^k is a vector $\mathbb{k}(p_{Bb}^h, p_{Aa}^k)$ of the form

$$\left[A, \left\lfloor \frac{t(p_{Bb}^h) - t(p_{Aa}^k)}{\delta t} \right\rfloor, \left\lfloor \frac{\alpha(p_{Bb}^h)}{\delta \alpha} \right\rfloor, \left\lfloor \frac{\alpha(p_{Aa}^k)}{\delta \alpha} \right\rfloor, \left\lfloor \frac{v(p_{Bb}^h) - v(p_{Aa}^k)}{\delta v} \right\rfloor \right],$$

where $\delta \alpha$, δt and δv are the units (in degrees, seconds, and meters/second, respectively) used to discretize angles, times and speeds. A value is instead a vector of four fields: (i) the number of times that the RSU observed a couple of production phases with characteristics as from the associated key; (ii) the number of times that vehicles from the aforementioned couples of production phases actually generated a forward phase; (iii) the average time elapsed between the start of the last production phase and the start of the forward phase, if any of the latter has ever occurred; (iv) the average duration of the forward phase, if any has ever occurred.

Contacts map store information on the probability and timing of contacts among vehicles, and thus implement a prediction model. By exchanging information about their respective local production phases, RSUs can leverage their contact maps so as to take scheduling decisions. More precisely, those decision are based on the computation, for each active local production phase, of the *delivery potential* \mathbb{P}_a resulting from electing one (or some, or all) of the local vehicles as relay(s) for data destined to a specific downloader vehicle a . The delivery potential is obtained as the sum of the individual contact probabilities \mathbb{P}_b , derived from assigning data for the downloader a to each elected local carrier b ¹⁰. The process is repeated for each known

¹⁰Note that, unlike the optimal and probabilistic graph-based schemes, the contact map-based solution exploits multiple relays for the same content chunk. This improves reliability but does not impact airtime utilization at the RSU, since, thanks to the broadcast nature of the wireless channel, a single transmission is sufficient to transfer the same data to all elected local relays.

downloader car, and, in the end, the downloader vehicle associated with the highest delivery potential \mathbb{p} is chosen as the target of a carry&forward transfer through local carriers that contributed to \mathbb{p} . Note that \mathbb{p} is a potential and not a probability: indeed, \mathbb{p} can be higher than one to counter uncertainties in probability estimates.

The framework for carriers selection run at a generic RSU B is shown as pseudocode in Fig. 3.10. There, priority is always given to direct data transfers to downloader cars, and fairness among them is provided by always picking the vehicle that is the closest to the RSU. The parameter \mathbb{P}_{min} controls the minimum delivery potential required to attempt cooperative download through local carriers. The value of such parameter (line 01 in Fig. 3.10), together with the way the delivery potential \mathbb{P}_{Aa}^k associated to the downloader production phase p_{Aa}^k and its relative carriers list \bar{c}_{Aa}^k are updated (lines 12 and 13 in Fig. 3.10), can be varied so as to define different relay selection algorithms. We evaluated four algorithms, as follows.

- The *Blind* relay selection algorithm aims at fully exploiting the airtime available at RSUs, by delivering data to all available local carriers whenever possible. This algorithm does not make use of the contacts map, but randomly chooses a downloader car as the destination of the data: we thus employ it as a benchmark for the other schemes.
- The *p-Driven* relay selection algorithm is a contact map-based version of the Blind algorithm. It still tries to exploit as much as possible the RSU wireless resources, but this time the RSU selects the target downloader that maximizes the sum of its contact probabilities with all the local carriers.
- The *p-Constrained* relay selection algorithm builds on top of the p-Driven scheme, and enforces that local vehicles with individual contact probability \mathbb{p}_b lower than $\mathbb{P}_{ind} > 0$ are not considered for relaying. Also, \mathbb{P}_{min} is set to a value higher than 0, so that downloader vehicles with delivery potential \mathbb{p}_a lower than \mathbb{P}_{min} are discarded.
- The *(p,t)-Constrained* relay selection algorithm adds time constraints to the probability bounds of the p-Constrained scheme. It exploits information on the average time to contact and contact duration to estimate the time interval during which the local vehicle will meet the downloader car, so as to improve the delivery precision and possibly reduce the data carriers involved in the cooperative download.

The four algorithms are ordered by increasing complexity and accuracy in the relay selection, leading to a better use of wireless resources. However, quality could come at cost of quantity, as discarding too many candidate relays may hinder potentially successful cooperation among vehicles.

We evaluated the performance of the contact map-based scheduler in several typical urban scenarios, in presence of 1 RSU/km² deployed at crowded intersections. Around 1% of vehicles concurrently request infinite-size contents. We compared the performance of the different relay selection algorithms to an *Oracle* scheme, which builds contact maps on perfect knowledge of the future trajectories of all vehicles, and thus approximates well the optimal solution.

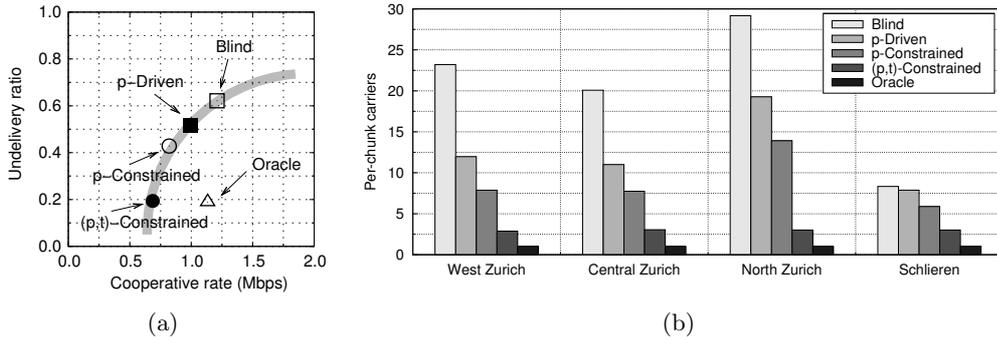


Figure 3.11: Cooperative download via a contact map-based scheduling. (a) Working points of the different relay selection algorithms in the space of download rate and undelivery ratio. (b) Number of relays carrying a same content chunk for different relay selection schemes, in different road traffic scenarios.

Fig. 3.11 shows the result of each algorithm, aggregated over all scenarios, based on two key metrics: (i) the download rate, on the y axis, is the average content transfer speed experienced by downloaders¹¹; (ii) the undelivered ratio, on the x axis, is the average ratio of content chunks that are not delivered to a downloader vehicle, over all those scheduled for that vehicle. The working points of the different techniques in the space of download rate and undelivery ratio evidence a tradeoff between the transfer volume and the reliability of the cooperative download process. In particular, the non-linear distribution of the working points, evidenced by the grey curves in the plot, seems to indicate that the additional complexity of the (p,t)-Constrained relay selection pays off, yielding an efficiency gain that compensates for the download rate reduction. The effectiveness of the (p,t)-Constrained scheduler is even clearer when looking at the mean number of relays carrying a same content chunk to a given target downloader, in Fig. 3.11(b). This plot separates the results obtained in four urban scenarios: in all cases, more precise algorithms result in a smaller number of relays per chunk, and thus in lower RSU loads. In other words, the higher precision in the scheduling granted by the (p,t)-Constrained scheme yields a significantly lower charge on the infrastructure and a reduced demand of resources by cooperating vehicles.

3.1.3 Discussion

In this Section, we considered a cooperative download system that uses I2V and V2V communication so as to transfer content from the Internet towards downloaders on-board vehicles. Such a system significantly differs from those targeting information dissemination or cooperative caching, since it addresses the more general case of users interested in different contents each. Therefore, a dedicated study is required,

¹¹Note that the absolute values of the download rate in these plots are not comparable with those of the optimal solution and of the probabilistic graph-based scheduling, due to a different (and much more conservative) network-level data transmission rate model. In any case, one can use the Oracle algorithm as a reference for quasi-optimal scheduling.

which was carried out by outlining the performance bounds of cooperative download as well as by devising practical solutions that implement that principle. The results obtained with optimal and practical approaches are in agreement and allow us to draw conclusions on several important properties of cooperative download in vehicular environments, as follows.

- There exist two operation regimes, characterized by different performance and impact of the system settings, that depend on the level of vehicular communication technology penetration. By considering the performance bounds in a typical urban scenario, the watershed arises when 20-30% of the vehicles participate in the network.
- The contribution to performance of V2V traffic relaying is critical. It can compensate for reduced coverage as well as for a non-optimal RSU placement, with such an effect becoming more and more evident as the technology penetration rate increases. The contribution of V2V communications remains relevant even under a pervasive RSU deployment and in both penetration regimes, as optimal scheduling tends to favor high-rate V2V transfers over low-rate I2V communications.
- Accurate anticipation of user mobility is paramount to the system performance, since most of the V2V traffic relaying takes place through the carry-and-forward paradigm. However, the complexity of multi-hop protocols can be limited to one relay, as the contribution of transfers over a higher number of hops is negligible. This confirms findings on bus networks [10, 11], which thus apply also to a more general vehicular downloading context.
- The structure of the road topology and the route followed by vehicles determine the downloading performance experienced by the users. Thus, one should adapt the system configuration to the characteristics of the road environment. As a consequence, unfairness among users should be expected, at least unless there is a combination of (i) pervasive presence of RSUs and relays, and (ii) limited number of concurrent downloaders.
- The accuracy of the mobility prediction has a relatively small impact on the actual download performance, but it significantly affects the system efficiency. Indeed, precise forecasting contacts among vehicles can spare wireless resources, better serve downloaders on secondary roads, reduce the download latency, and lower the RSU deployment cost.

In conclusion, our results show that cooperative download via I2V and V2V communication can represent a viable alternative, or a complementary solution, to traditional via cellular networks for mobile users willing to retrieve contents from the Internet.

3.2 The uplink direction: floating car data collection

Uplink data transfers from vehicles to the Internet can be mapped to the collection of so-called floating car data (FCD). These consist of information generated by moving vehicles and uploaded to Internet-based control centers for processing and analysis.

FCD is today employed for, e.g., distant monitoring of on-board Electronic Control Units (ECUs). ECUs, whose number varies between 30 for low-end cars and 100 for premium-class automobiles [12], locally oversee in-vehicle operations, controlling almost all car functionalities. Systems such as BMW Assist, Ford SYNC, General Motor OnStar, Toyota Safety Connect and Mercedes-Benz mbrace, just to cite a few representative examples, retrieve the data generated by the ECUs in the form of FCD, so as to provide seamless distant support to the driver and the passengers. Services cover safety, diagnostic and anti-theft applications. Another common practical use of FCD is in the field of real-time road traffic monitoring. As an example, technologies such as TomTom HD Traffic and Meihui TrafficCast leverage FCD carrying anonymized vehicle position and speed so as to determine the traffic conditions in real-time and offer more efficient navigation services.

The FCD upload is today performed individually by each participating vehicle via the cellular infrastructure, as shown in Fig. 3.12(a). Since the present negligible penetration rates of FCD-based technologies allow the cellular infrastructure to accommodate the FCD uplink traffic, completely relying on the pervasive access network is a convenient practice. However, the success of the few FCD-based services that have been deployed is fostering a large-scale adoption of FCD-based solutions. As a significant example, FCD is going to play a key role in pervasive urban sensing: vehicles would collect environmental information about the metropolitan areas they travel through, and upload such data to Internet-based control centers for fusion and analysis. Urban sensing is indeed envisioned to significantly improve our understanding of urban dynamics and is commonly regarded as a fundamental component in forthcoming smart cities. Additionally, many of the existing and future usages of FCD require the harvesting of data from the largest possible vehicle population. For example, for the distant monitoring of ECUs, each single car must be continuously probed, while in real-time traffic monitoring or in urban sensing, the quality of the aggregate information significantly improves with the number of available (positioning or sensing) samples. Not only the quantity, but also the frequency of the FCD collection is foreseen to progressively increase, driven by the need for growing accuracy in the monitoring or sensing activities.

These observations let us reasonably speculate that the consolidation of FCD-based technologies reaching near-100% penetration rates risks to induce a non-negligible load on the cellular uplink access. Considering that the mobile demand has reached the capacity limits of 3G networks [3], and that even the upcoming LTE infrastructure is already deemed unable to cover such a gap [2], offloading FCD could only benefit the cellular network operation. To that end, direct vehicle-to-vehicle (V2V) communication based on Dedicated Short Range Communication (DSRC) could come in handy. The standardization activity has recently led to a number of protocol stack proposals, including IEEE 802.11p, IEEE 1609.x, ETSI ITS G5 and ISO CALM, deemed to enable communication among vehicles traveling within a range of a few hundred meters. Local gathering and fusion via V2V transfers could be leveraged for FCD offload as shown in Fig. 3.12(b). There, a subset of the vehicles gather the data sensed by neighboring cars through DSRC communication, and fusion it with their own observations before uploading the aggregate information via

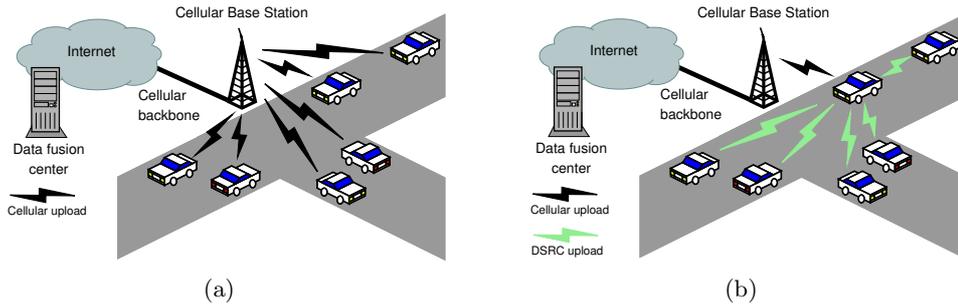


Figure 3.12: FCD upload scenarios. (a) Traditional via the cellular infrastructure. (b) Offloaded through V2V transfers.

a single cellular network transfer. The system would result in significant offload of the cellular infrastructure, in terms of channel signalization, control header overhead and, depending on the local fusion level, sheer uplink traffic volume.

In this Section, we explore the above a scenario for offloading the cellular infrastructure from FCD uploads. We remark that the FCD offload problem is different from those of downlink or WiFi-based offloading, and requires a dedicated analysis. More precisely, offloading FCD uploads maps to: (i) identifying the subset of vehicles in charge of performing the data fusion and upload, so as to harvest the maximum FCD amount; (ii) doing so in an efficient, distributed way.

3.2.1 Drawing the performance bounds

We first approach the problem from an oracle viewpoint, and formulate FCD offloading as an optimization problem. The solution of such a problem allows us to outline the performance bounds of the system.

Our goal is then to identify a set of vehicles, each of which gathers FCD from its communicating neighbors, and performs the data fusion and upload. Clearly, we wish such a set to be as small as possible, since the fewer the vehicles performing the local fusion and upload, the lower the uplink traffic load on the cellular network. At the same time, however, we do not want the offload process to reduce the quality of the overall FCD information: this maps to the requirement that the V2V gathering of Fig. 3.12(b) should collect the same FCD that would have been individually uploaded by cars in the infrastructure-based approach of Fig. 3.12(a). In other words, the objective becomes that of collecting FCD from the whole network using the least amount of vehicles, so as to minimize the number of uploads and maximize the local FCD fusion. The FCD offloading problem can thus be formulated as a Minimum Dominating Set (MDS) problem, whose solution yields the minimal set of vehicles that cover all the other cars through V2V communication.

The problem of finding a MDS is typically NP-hard. However, several heuristics exist that can compute approximate solutions in an efficient way. Specifically, we consider the greedy algorithm proposed by Lovasz [13], a simplified version of the same that exploits the related problem of the Maximum Independent Set (MIS) and is named Lexicographically First MIS (LFMIS) [14], and solutions by Marathe *et*

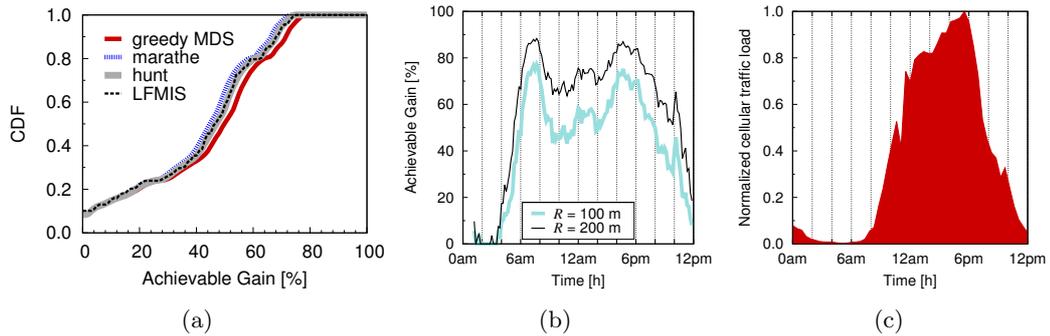


Figure 3.13: Offloaded FCD upload performance bounds. (a) CDF of the optimal system gain over 24 hours. (b) Dynamics of the optimal system gain over time. (c) Typical normalized traffic load at the cellular access network, over a 24-hour period. The latter data was obtained from experiments conducted by the Autonomous Networks Research Group at USC, <http://anrg.usc.edu>.

al. [15] and Hunt *et al.* [16] that leverage properties of unit disk graphs to reduce the approximation with respect to the optimal MDS solution.

We apply the different MDS approximation algorithms to the Cologne road traffic dataset presented in Sec. 2.1. Fig. 3.13(a) shows the Cumulative Distribution Function (CDF) of the system gain, i.e., the fraction of vehicles that do not have to access the cellular infrastructure when FCD is offloaded through DSRC communication¹², at each time instant of the day. We can observe that the system gain can be significant: more than 70% of the time we are theoretically able to offload more than 40% of the FCD through V2V communication. In some situations the gain can even reach 80% of the offered load. Also, there is no appreciable difference among the four heuristics, which lead to relatively close results.

The shape of the gain distribution demands for an investigation of when and where the offload is more or less effective. The evolution of the system gain over time is portrayed in Fig. 3.13(b). We can note that the gain varies significantly depending on the hour considered: in fact, its dynamics closely follow those of the road traffic activity, in Fig. 2.1(c). Indeed, when the number of cars is very low, e.g., before 5 am or after 11 pm, most vehicles are isolated from a V2V communication viewpoint, and thus forced to upload their own FCD individually. However, we remark that this would not be a major problem, as the number of implicated vehicles is extremely low, and the cellular infrastructure is almost not utilized during those time periods, as shown by Fig. 3.13(c), which portrays a typical daily load pattern of a cellular base station. More interestingly, V2V communication can offload 50% to 60% of FCD during non-rush traffic hours, e.g., between 9 am and 4 pm. Since that period yields moderate-to-high data traffic load in Fig. 3.13(c), FCD offload could have a significant impact there. Finally, it is during the traffic peak periods, i.e., between

¹²Such an intuitive metric has the advantage of being of general validity throughout different cellular technologies and FCD applications, which all map into scaling factors to the gain. See the paper appeared at IEEE WoWMoM 2013 for further details.

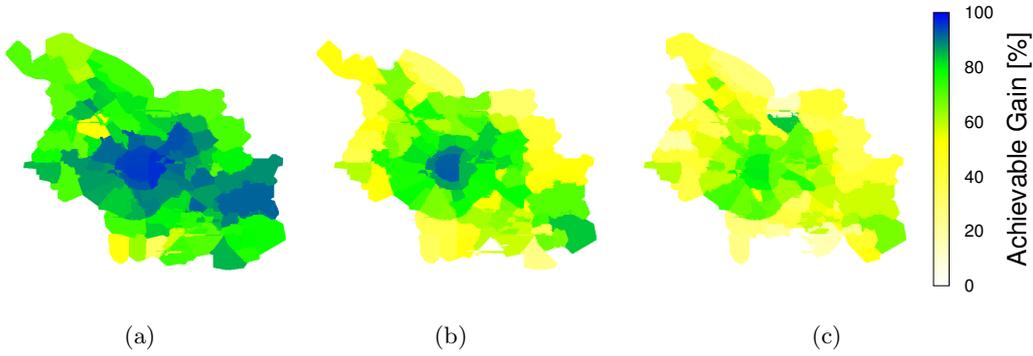


Figure 3.14: Spatial dynamics of offloaded FCD upload, obtained by separating the gain on the 86 *Stadtteile*, i.e., districts, of Cologne. (a) System gain at 7.30 am. (b) System gain at 12.30 am. (c) System gain at 9 pm. Figure best viewed in colors.

7 am and 9 am and between 4 pm and 6 pm, that the DSRC-based FCD offloading attains its maximum gain, at around 75%. While the first time period is a low-load one for the cellular infrastructure, the second maps exactly to the daily peak in the cellular data traffic load. We can thus conclude that during several hours in the afternoon high-gain FCD offload may prove paramount to reduce the load on the cellular infrastructure.

Not only temporal dynamics, but also spatial ones are of interest to our study. We analyze the impact of FCD offloading from a geographical perspective in Fig. 3.14, where the plots portray the gain in each of the 86 Cologne *Stadtteile*, i.e., districts, at three different hours. The distribution of the gain is far from being uniform over the region. The city center is the area with the highest gain, reaching values around 0.95 at 7.30 am, which means that 95% of the vehicles can avoid accessing the cellular network and still contribute to the FCD collection. The South-East suburban area, where Porz, the largest borough of Cologne, is situated, is also characterized by consistently higher gain than most other areas. Other districts, such as the southern ones, show instead low FCD offloading gain over the whole day. Once more, the gain is largely dependent on the road traffic volume. The downtown or Porz areas are where most of the vehicular activity gathers, which results in higher vertex degrees of the V2V connectivity graph and higher margin of gain through the MDS uploader selection. Most importantly, these regions are again those characterized by the highest human activity, and thus by the highest cellular network load. Therefore, these are also where the FCD offloading will be more critical.

Overall, our results showed that FCD offloading through V2V communication is an interesting approach to reduce the volume of traffic that will be uploaded to the cellular infrastructure by sensing vehicles in urban areas, and this DSRC-based FCD offloading performs best precisely at the hours and in the regions where the cellular network needs to be relieved the most.

3.2.2 Practical solutions

Having identified the significant gain attainable through DSRC-based FCD offloading, we are now interested in defining practical solutions that exploit such a networking paradigm. This implies relaxing the oracle assumption and designing distributed solutions to the MDS problem in a dynamic vehicular network. We thus proposed and analyzed two heuristics that allow the distributed construction of a set of relays used to offload FCD in a short timeframe, so as to adapt to the fast dynamics of V2V communications.

Degree-Based with Confirmation (DB-C)

The first mechanism is named Degree-Based with Confirmation (DB-C). It leverages the fact that vehicles using DSRC natively obtain information regarding their neighbors by the means of beacons periodically transmitted on the so-called control channel, as dictated by ETSI (via the Cooperative Awareness Messages, or CAM) and by the SAE J2735 dictionary set. Such knowledge can be used to distributely compute the degree of each node in the vehicular network, an information then used to decide whether a vehicle shall act as an uploader or not. Specifically, the DB-C scheme runs in two phases.

- In the first phase, each vehicle i with D_i neighbors (including vehicle i itself, therefore $D_i \geq 1$) uploads its data as well as that of its V2V neighbors on the cellular link with a probability k/D_i , upper-bounded at 1. There, k is a system parameter that will be discussed next.
- In the second phase, we ensure that a dominating set is built, which cannot be guaranteed by the probabilistic approach of the first phase. To that end, a simple confirmation mechanism is used: each vehicle choosing during the first phase to act as uploader for its neighbors informs them on the V2V channel. This allows every node to learn if its data was uploaded in the first phase or not. In the latter case, the vehicle can upload its own information via the cellular network, and become part of the dominating set.

The DB-C scheme relies on the parameter k , which regulates the size of the dominating set. As a matter of fact, a too low k leads to many vehicles left uncovered during the first phase and thus resorting to the traditional cellular upload during the second phase; conversely, a too high k results in an exceeding number of vehicles self-electing into the dominating set during the first phase, reducing the system gain. In order to determine the optimal value of k , we formulate an optimization problem, as follows.

First, we express the number of vehicles that are covered or not covered at the end of the first phase, as a function of k . The probability that a node i with $D_i = d$ neighbors is not covered during the first phase depends on $P_{nt}(d)$, i.e., the probability that the node itself does not transmit, and $P_{nn}(d)$, i.e., the probability that all its one-hop neighbors do not act as uploaders:

$$P_{nc}(d) = P_{nt}(d) \cdot P_{nn}(d). \quad (3.4)$$

The first probability is obtained directly from the definition of the DB-C scheme operation:

$$P_{nt}(d) = 1 - \min(1, k/d).$$

The second probability can be instead written as:

$$P_{nn}(d) = \prod_{j:v_j \in \mathbb{V}_i^1} (1 - \min(1, k/D_j)),$$

where \mathbb{V}_i^1 is the set containing all the one-hop neighbors of node i . From the discussion in Sec. 2.2.3, vehicular networks are highly assortative, thus the neighbors of a vehicle can be safely assumed to have degrees similar to that of the vehicle itself. Thus, $D_j \approx D_i = d$ and, by replacing the terms in (3.4), we obtain:

$$P_{nc}(d) \approx (1 - \min(1, k/d))^d. \quad (3.5)$$

We then assume that (i) traditional upload through the cellular infrastructure has a cost C_c for each vehicle, (ii) aggregate upload by one vehicle acting as a relay for d neighbors has a cost of $C_v(d)$ for the uploader, and (iii) V2V communication has no cost. Then, the average transmission cost of one FCD collection period can be expressed as

$$C = \sum_{d=1}^{\infty} \pi_d [C_c \cdot P_{nc}(d) + C_v(d) \cdot (1 - P_{nc}(d))]. \quad (3.6)$$

Without loss of generality, we can assume that the cost in each scenario is given by the number of transmissions on the uplink, namely $C_c = 1$, and $C_v(d) = \min(1, k/d)$. Replacing the terms and using the result in (3.5), (3.6) becomes:

$$C = \sum_{d=1}^{\infty} \pi_d \left\{ 1 \cdot \left(1 - \min\left(1, \frac{k}{d}\right)\right)^d + \min\left(1, \frac{k}{d}\right) \left[1 - \left(1 - \min\left(1, \frac{k}{d}\right)\right)^d\right] \right\}. \quad (3.7)$$

We can separate the sums in (3.7) according to whether the node degree is below or above the parameter k , and obtain:

$$C = \sum_{d=1}^k \pi_d \frac{d}{d} + \sum_{d=k+1}^{\infty} \pi_d \left\{ \left(1 - \frac{k}{d}\right)^d + \frac{k}{d} \left[1 - \left(1 - \frac{k}{d}\right)^d\right] \right\}. \quad (3.8)$$

After some simple computations, the average transmission cost is:

$$C = \sum_{d=1}^k \pi_d + \sum_{d=k+1}^{\infty} \left[\pi_d \frac{k}{d} + \pi_d \left(1 - \frac{k}{d}\right)^{d+1} \right]. \quad (3.9)$$

In (3.9), the first sum is the cost brought by those nodes which are sure to be selected by the DB mechanism because they have $k-1$ or less neighbors. The second term, of the second sum is the cost of those transmitting on the cellular uplink using the DB-C approach, as they remained uncovered by DB. Finally, the first term of the second sum gives us a gain over the traditional scenario, and represents the cost of

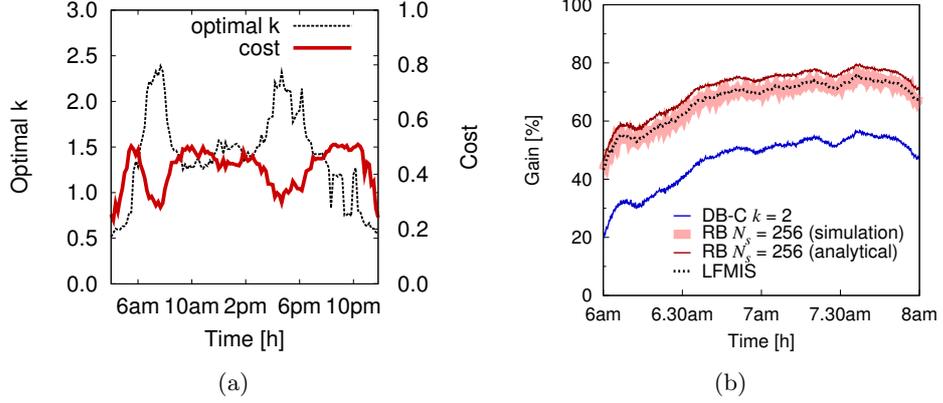


Figure 3.15: Practical schemes for offloaded FCD upload. (a) Optimal k and minimal cost of PR-C. (b) Performance of the PR-C and RB mechanisms during the morning road traffic peak.

those nodes relaying FCD for their neighbors in the first transmission round. The optimal value of k is thus obtained by solving:

$$\min \sum_{d=1}^k \pi_d + \sum_{d=k+1}^{\infty} \left[\pi_d \frac{k}{d} + \pi_d \left(1 - \frac{k}{d}\right)^{d+1} \right], \quad (3.10)$$

where k is a continuous variable. The problem is not trivial to solve efficiently, as the variable k appears an unbounded (a priori) number of times, with different signs and exponents. The first fact, i.e., the $(1 - \frac{k}{d})^{d+1}$ term, makes it quite complex (although not entirely impossible) to analytically derive the expression and determine its minimum. The second fact, i.e., that k happens to appear with a negative coefficient, rules out other straightforward optimization techniques such as convex optimization of posynomial functions [17]. On the other hand, the expression to be minimized has three desirable properties: (i) it is univariate, i.e., includes a single variable k ; (ii) it is continuous and derivable with respect to k ; (iii) the values of k can be bounded, e.g., between 0 and D_{max} . Root-finding and optimization of univariate expressions has been long studied, and there are plenty of existing solutions to the problem – most of which are similar, in principle, to the well-known bisection method. *Well-behaved* functions, especially derivable ones, are associated to faster convergence times. Finally, if the optimum can be searched for in a limited interval, the result can be guaranteed to be a global optimum (as opposed to a local one). In our case, we solve (3.10) via the Brent method [18, Chapter 7], which runs in polynomial (namely, quadratic) time and returns a global optimum.

The optimization results are presented in Fig. 3.15(a). The optimal value of k is generally between 1.5 and 2.5, and it follows the evolution of the vehicular density, with two peaks at around 7:30 am and 6:00 pm. Also, the cost achieved by this optimal k is smaller when the number of vehicles increases, as more information can be offloaded via DSRC-based V2V communication.

Reservation-Based (RB)

Although DB-C results in a dominating set, its performance can still be improved when the vehicular density increases. In order to reduce this gap, we propose a simple reservation mechanism which selects uploaders in a single V2V communication round. As discussed below, this mechanism results in a constant approximation of the MDS in an ideal scenario and in small (although not minimal) dominating sets in a realistic DSRC-based network.

This mechanism, named Reservation-based (RB), runs through the following steps at each FCD collection period:

- At the beginning of every collection period, there is a reservation phase (assumed slotted and containing N_s slots), where each vehicle selects a transmission slot among the N_s available and enters the *contender* state where it backs-off, waiting for the chosen slot.
- In every time slot, the vehicles in the *contender* state having chosen the corresponding back-off transmit a reservation message and change their status to *dominator*.
- A vehicle in the *contender* state and receiving a reservation message from one of its neighbors becomes *dominated* and cancels its back-off.

An overview of the performance of the DB-C and RB schemes is provided in Fig. 3.15(b). There, we focus on 2 hours of road traffic, between 6:00 am and 8:00 am. The plot reports the system gain recorded under the optimal solution (LFMIS), the DB-C and the RB schemes. As far as DB-C is concerned, we used $k = 2$, a constant value that well approximates those suggested by the solution of the optimization problem in (3.10) and shown in Fig. 3.15(a). Two curves are instead presented for the RB scheme, which display the gain attained in simulation and with an analytical model of the mechanism¹³, when $N_s = 256$. We remark that the DB-C is able to reach a gain of around 50% at peak hours, however, it also shows a somehow constant gap of 25% system gain from the optimal solution. Instead, the gain attained with the RB mechanism is not distinguishable from the one produced by the LFMIS. As a result, RB allows to spare up to 80% of uplink transmission via the cellular infrastructure at 7:30 am.

3.3 Conclusions of the Chapter

In this Chapter, I presented results relevant to data transfers in environments that involve communication-enabled vehicles. Among the major contributions of the research carried out in that context, one can list: (i) the definition of the performance bounds of data transfers through DSRC-based infrastructure-to-vehicle and vehicle-to-vehicle communication, both in the downlink and uplink directions; (ii) the demonstration that communication among vehicles can help in a significant manner the download and upload of data by users onboard them, with peaks of 80% of

¹³The analytical model is not introduced here for the sake of brevity. Please refer to the paper appeared at IEEE WoWMoM 2013 for a complete discussion.

the traffic managed through DSRC only (and thus offloaded from the cellular infrastructure) in both cases. (iii) the provision of general design guidelines as well as practical solutions that efficiently address (large content) downlink and (floating car data) uplink transfers, achieving results that are close to the optimal ones.

Bibliography

- [1] Ericsson, “*Mobile Data Traffic Surpasses Voice*,” Press Release, March 2010.
- [2] Cisco, “*Global Mobile Data Traffic Forecast Update, 2012-2017*,” White Paper, February 2013.
- [3] U. Paul, A.P. Subramanian, M.M. Buddhikot, S.R. Das, “*Understanding traffic dynamics in cellular data networks*,” IEEE INFOCOM, Shanghai, China, April 2011.
- [4] R. Pries, F. Wamser, D. Staehle, K. Heck, P. Tran-Gia, “*Traffic measurement and analysis of a broadband wireless Internet access*,” IEEE VTC-Spring, Barcelona, April 2009.
- [5] F. Aidouni, M. Latapy, C. Magnien, “*Ten weeks in the life of an eDonkey server*,” Hot-P2P, Rome, Italy, May 2009.
- [6] L. R. Ford, D. R. Fulkerson, “*Flows in networks*,” Princeton University Press, Princeton, NJ, USA, 1962.
- [7] D. Hadaller, S. Keshav, T. Brecht, S. Agarwal, “*Vehicular opportunistic communication under the microscope*,” ACM MobySys, San Juan, Puerto Rico, June 2007.
- [8] J. Mattingley, S. Boyd, “*Real-time convex optimization in signal processing*,” IEEE Signal Processing Magazine, 27(3), 2010.
- [9] J. Ott, D. Kutscher, “*Drive-thru Internet: IEEE 802.11b for automobile users*,” IEEE INFOCOM, Hong Kong, China, March 2004.
- [10] A. Balasubramanian, B.N. Levine, A. Venkataramani, “*Enhancing interactive Web applications in hybrid networks*,” ACM/IEEE MobiCom, San Francisco, CA, September 2008.
- [11] N. Banerjee, M.D. Corner, D. Towsley, B.N. Levine, “*Relays, meshes, base stations: Enhancing mobile networks with infrastructure*,” ACM/IEEE MobiCom, San Francisco, CA, September 2008.
- [12] R.N. Charette, “*This car runs on code*,” IEEE Spectrum, February 2009.
- [13] L. Lovasz, “*On the Ratio of Optimal Integral and Fractional Covers*,” Discrete Mathematics, 13(4):383–390, December 1975.

- [14] M. Luby, “A Simple Parallel Algorithm for the Maximal Independent Set Problem,” ACM STOC, Providence, RI, USA, May 1985.
- [15] M. Marathe, H. Bredt, H. Hunt III, S. Ravi, D. Rosenkrantz, “Simple Heuristics for Unit Disk Graphs,” *Networks*, 25(2):59–68, March 1995.
- [16] H. Hunt III, M. Marathe, V. Radhakrishnan, S. Ravi, D. Rosenkrantz, R. Stearns, “NC-Approximation Schemes for NP- and PSPACE-Hard Problems for Geometric Graphs,” *Journal of Algorithms*, 26(2):238–274, February 1998.
- [17] S.P. Boyd, L. Vandenberghe, “*Convex Optimization*,” Cambridge University Press, 2004.
- [18] R.P. Brent, “*Algorithms for Minimization Without Derivatives*,” Dover Publications, 2002.

Chapter 4

Secure positioning in vehicular environments

Location awareness is an asset in vehicular networks, where a wide range of protocols and applications require knowledge of the position of the moving vehicles. Geographic routing, data gathering, movement coordination, location-specific services, danger warning, traffic monitoring, and driver liability are all examples of services that build on the availability of vehicle positioning information. The latter are typically generated and advertised by the vehicles themselves: the verification of the correctness of location data becomes then an all-important issue, even more so in the presence of malfunctioning or adversarial users in the system. In these cases, we need solutions that make it possible to validate the positions announced by vehicles, so as to eventually detect and evict users announcing false locations.

This Chapter presents solutions to the problem of securing the positioning information advertised by communication-enabled vehicles. I will first introduce a distributed solution intended for individual vehicles willing to verify the location of their neighbors, in Sec. 4.1. Then, I will describe a centralized framework, where a Location Authority is in charge of assessing the trustworthiness of the position announced by networked vehicles, in Sec. 4.2.

The solutions outlined next are the result of collaborations with Francesco Malandrino, Carlo Borgiattino, Roberto Sadao, Carla-Fabiana Chiasserini, Claudio Casetti, and Panos Papadimitratos. The related publications are as follow:

- M. Fiore, C. Casetti, C.-F. Chiasserini, P. Papadimitratos, “*Secure Neighbor Position Discovery in Vehicular Networks*,” IEEE/IFIP MedHocNet 2011, Favignana, Italy, June 2011.
- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, R.S. Yokoyama, C. Borgiattino, “*A-VIP: Anonymous Verification and Inference of Positions in Vehicular Networks*,” IEEE INFOCOM 2013 Miniconference, Turin, Italy, April 2013.
- M. Fiore, C. Casetti, C.-F. Chiasserini, P. Papadimitratos, “*Discovery and Verification of Neighbor Positions in Mobile Ad Hoc Networks*,” IEEE Transactions on Mobile Computing, 12(2), January 2013.

- F. Malandrino, C. Borgiattino, C. Casetti, C.-F. Chiasserini, M. Fiore, R.S. Yokoyama, “Verification and Inference of Positions in Vehicular Networks through Anonymous Beacons,” IEEE Transactions on Mobile Computing, to appear.

4.1 Distributed secure positioning

We first consider the case of a pure ad hoc network paradigm, where a pervasive infrastructure is not present, and the location data must be obtained through vehicle-to-vehicle communication. Such a scenario leaves the door open for adversarial users to misuse or disrupt the location-based services. For example, by advertising forged positions, adversaries could bias geographic routing or data gathering processes, attracting network traffic and then eavesdropping or discarding it. Similarly, counterfeit positions could grant adversaries unauthorized access to location-dependent services, let vehicles forfeit road tolls, disrupt vehicular traffic or endanger passengers and drivers.

In this context, the challenge is to perform, in absence of trusted nodes, a fully-distributed, lightweight neighbor position verification (NPV) procedure that enables each node, i.e., vehicle, to acquire the locations advertised by its neighbors on demand, and assess their truthfulness. We therefore propose an NPV protocol that has the following features:

- It is designed for spontaneous ad hoc environments, and, as such, it does not rely on the presence of a trusted infrastructure or of a-priori trustworthy nodes;
- It leverages cooperation but allows a node to perform all verification procedures autonomously. This approach has no need for lengthy interactions, e.g., to reach a consensus among multiple nodes, making our scheme suitable for high-mobility vehicular environments;
- It is reactive, meaning that it can be executed by any node, at any point in time, without prior knowledge of the neighborhood;
- It is robust against independent and colluding adversaries;
- It is lightweight, as it generates low overhead traffic.

Additionally, our NPV scheme is compatible with state-of-the-art security architectures, including the ones that have been proposed for vehicular networks [1, 2], which represent a likely deployment environment for NPV.

4.1.1 NPV protocol operation

We propose a fully-distributed cooperative scheme for NPV, which enables a node, hereinafter called the *verifier*, to discover and verify the position of its communication neighbors. For clarity, here we summarize the principles of the protocol as well as the gist of its resilience analysis¹.

¹Detailed discussions of message format, verification tests and protocol resilience are provided in the paper appeared in IEEE Transactions on Mobile Computing 12(2).

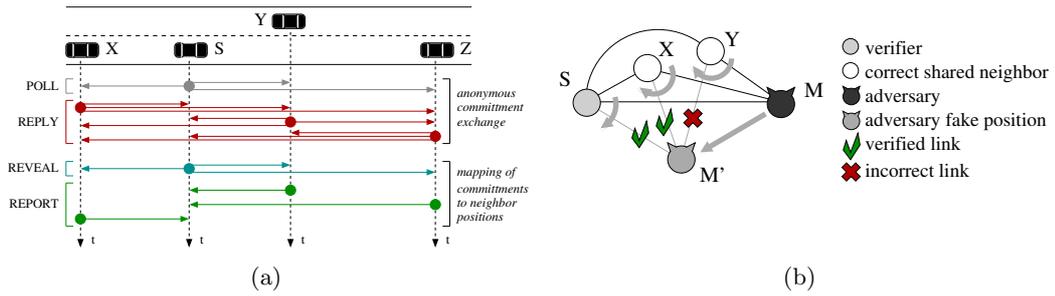


Figure 4.1: Neighbor position verification (NPV) protocol operation. (a) Message exchange overview during one instance of the protocol. (b) Example of topological information stored by verifier S at the end of the message exchange and effect of a fake position announcement by M .

A verifier, S , can initiate the protocol at any time instant, by triggering the 4-step message exchange depicted in Fig. 4.1(a), within its one-hop neighborhood. The aim of the message exchange is to let S collect information it can use to compute distances via time-of-flight (ToF)-based radio-frequency ranging², between any pair of its communication neighbors. To that end, POLL and REPLY messages are first broadcasted by S and its neighbors, respectively. These messages are anonymous and take advantage of the broadcast nature of the wireless medium, allowing nodes to record reciprocal timing information without disclosing their identities. Then, after a REVEAL broadcast by the verifier, nodes disclose to S , through secure and authenticated REPORT messages, their identities as well as the anonymous timing information they collected. The verifier S uses such data to match timings and identities; then, it uses the timings to perform ToF-based ranging and compute distances between all pairs of communicating nodes in its neighborhood.

Once S has derived such distances, it runs several position verification tests in order to classify each candidate neighbor as either:

1. *Verified*, i.e., a node the verifier deems to be at the claimed position;
2. *Faulty*, i.e., a node the verifier deems to have announced an incorrect position;
3. *Unverifiable*, i.e., a node the verifier cannot prove to be either correct or faulty, due to insufficient information.

Clearly, the verification tests aim at avoiding false negatives (i.e., adversaries announcing fake positions that are deemed verified) and false positives (i.e., correct nodes whose positions are deemed faulty), as well as at minimizing the number of unverifiable nodes. We remark that our NPV scheme does not target the creation of a consistent “map” of neighborhood relations throughout an ephemeral network: rather, it allows the verifier to independently classify its neighbors.

The basic principle the verification tests build upon is best explained by means of the example in Fig. 4.1(b). There, M is a malicious node announcing a false location

²ToF-based ranging at radio frequencies uses information on the transmission and reception times of a same message in order to estimate the distance between the message source and receiver. As discussed in [3], it is a reasonable assumption, although it requires modifications to off-the-shelf radio interfaces. Promising techniques for precise ToF-based RF ranging are currently marketed [4].

M' , so as to fraudulently gain some advantage over other nodes. The figure portrays the actual network topology (limited to M 's links for clarity of representation) with black edges, while the modified topology, induced by the fake position announced by M , is shown with gray edges. It is evident that the displacement of M to M' causes its edges with the other nodes to rotate, which, in turn, forces edge lengths to change as well. The tests thus look for discrepancies in the node distance information to identify incorrect node positions.

A malicious node, knowing the protocol, can try to outsmart the tests in a number of different ways. Overall, our NPV protocol guarantees that:

- An unknowledgeable adversary has no possibility of success against our NPV protocol;
- An independent knowledgeable adversary M can *move* at most two links (with the verifier S and with a shared neighbor X) without being detected: however, any additional link (e.g., with another shared neighbor Y) leads to inconsistencies between distances and positions that allow to identify the attacker: this is the situation depicted in Fig. 4.1(b). In a nutshell, independent adversaries, although knowledgeable, cannot harm the system;
- Colluding knowledgeable adversaries can announce timing information that reciprocally validate their distances, and pose a more dangerous threat to the system. However, we prove that an overwhelming presence of colluders in the verifier neighborhood is required for an attack to be successful.

4.1.2 NPV protocol performance evaluation

We evaluated the performance of our NPV protocol in a typical urban scenario, against adversaries that announce fake positions and aim at making the verifier validate them. We consider a worst-case attack model, with adversarial users that: (i) are internal, i.e., possess the cryptographic material to participate in the NPV protocol and try to exploit it, by advertising arbitrarily erroneous own positions or inject misleading information; (ii) know at each time instant the positions and identities of all their communication neighbors; (iii) run different attacks that are specifically designed against the NPV protocol; (iv) collude in doing so. Unless otherwise specified, adversaries amount to 5% of the total vehicles and are divided into groups of 5 colluders each. The results are shown in terms of the probability that the NPV protocol returns false positives and false negatives as well as of the probability that a (correct or adversary) vehicle is tagged as unverifiable. In all plot legends, C stands for correct vehicles, while M/Bas , M/Hyp and M/Col stand for adversaries mounting different types of colluding attacks.

Fig. 4.2(a) shows the robustness of the proposed NPV protocol to the density of adversaries, ranging from 5% to an overwhelming 30% of the total road traffic. Despite the massive presence of attackers, the probability that adversaries are verified (false negatives) increases ever so slightly with their density. A more evident effect is recorded on the probability of correct nodes being tagged as faulty (false positives), which however is a less harmful effect. Moreover, critical values – above 0.05 – of the probability of such false positives are only attained when at least one vehicle every

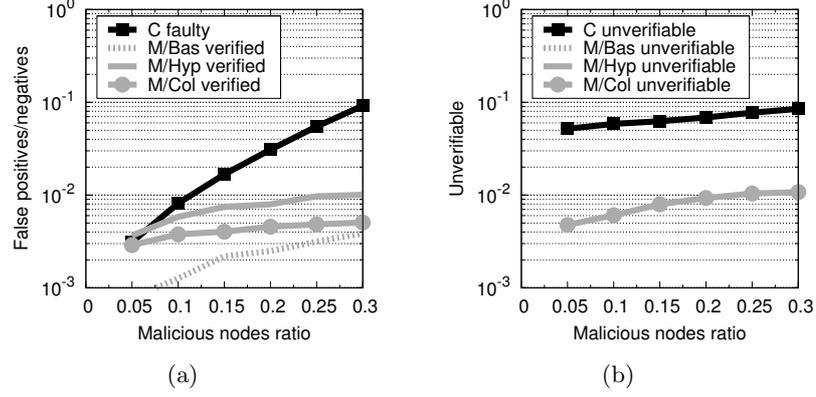


Figure 4.2: Neighbor position verification (NPV) protocol performance. Probability of incorrect classification versus the ratio of vehicles that are adversarial. (a) False negatives and false positives. (b) Unverifiables.

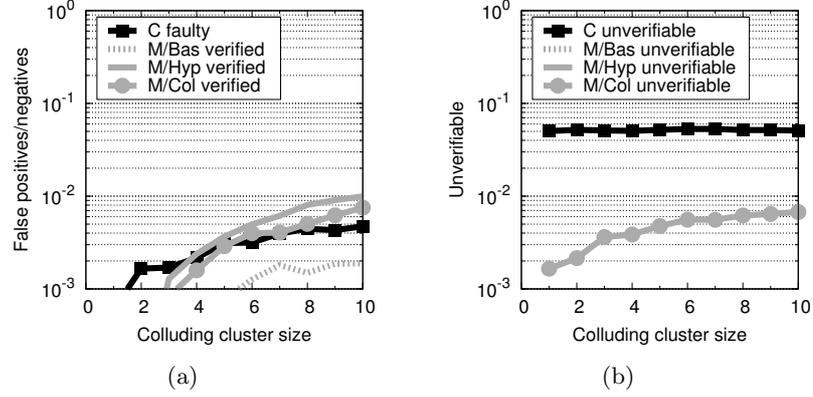


Figure 4.3: Neighbor position verification (NPV) protocol performance. Probability of incorrect classification versus the number of adversaries that collude in running each attack. (a) False negatives and false positives. (b) Unverifiables.

four is an adversary. As far as unverifiable nodes are concerned, their percentage, in Fig. 4.2(b), is always less than 1% for adversarial vehicles and between 5% and 9% for correct vehicles. We can remark how such values modestly increase with the growing presence of adversaries.

Not only the sheer number of attackers, but also their level of cooperation is an important parameter to account for. We thus examine the NPV protocol performance for different numbers of adversaries colluding in launching their attacks. The false negative and false positive probability in Fig. 4.3(a) clearly shows that an attack by colluder vehicles is most effective when carried out by at least three of them. In any case, the chance of wrong classification reaches 0.01 only for very large groups of cooperating adversaries, namely 10. The collusion level also affects the attackers' ability to disrupt the positioning of correct vehicles, which exhibit as high as a 0.4% chance to be tagged as faulty. Conversely, as shown in Fig. 4.3(b), the

collusion level does not cause more correct nodes to be unverifiable, since the main reason for correct nodes to be tagged as unverifiable is the lack of neighbors that can verify them. The chance for an adversary to be unverifiable increases with the colluder group size, although it is significant only for some types of attacks.

Overall, these and more tests³ showed that the proposed NPV protocol is highly resilient to attacks, even assuming excellent conditions for the adversaries to lead them. We observed typical probabilities of false positives and false negatives below 1%, while that of a node being tagged as unverifiable was below 5%. Moreover, additional results proved that a significant portion of the successful attacks yields small advantage to the adversaries in terms of displacement, and that the overhead introduced by the NPV does not exceed a few tens of kbytes even in the most critical conditions.

4.2 Centralized secure positioning

A different approach to the secure positioning problem is the centralized one. There, a Location Authority (LA) is in charge of collecting the position information periodically broadcasted by vehicles, and of verifying its validity. Such a scenario maps to that outlined by emerging standards on vehicular networking, such as ETSI ITS and the SAE J2735 dictionary set, which commend the transmission of periodic beacons by communication-enabled vehicles.

In that case, ensuring secure positioning must cope with three major problems, concerning (i) users' privacy, (ii) computational costs of security, and (iii) the system trust on user correctness. As for the first aspect, when not strictly required, public disclosure of the vehicle identity to all receiving devices in the proximity of a beaconer is an issue. Vehicles can be tracked, jeopardizing drivers' privacy and requiring complex pseudonym management [5]. Thus, there is a need for separating secure position identification by authorities and the possibility of undesirable user tracking by peers in the vehicular network. As for the second aspect, standard security mechanisms based on, e.g., asymmetric cryptography, induce significant protocol overhead and computational complexity. In fact, their use is recommended to be largely dependent on the applications and circumstances, and avoided whenever possible [6]. Finally, basic solutions (including direct upload of position information to the LA, via the cellular network) cannot guarantee the correctness of the location information provided by a user who owns the required cryptographic material, but has a malfunctioning GPS receiver or maliciously tampers with its GPS data.

In order to address the issues above, we proposed a framework for the Anonymous Verification and Inference of Positions (A-VIP). Unlike other solutions proposed in the literature, A-VIP features the following properties:

- it allows a trusted Location Authority (LA) to securely collect and verify the positions claimed by vehicles without resorting to computationally expensive asymmetric cryptography – as is instead done in the IEEE 1609.2 standard [1];

³Please refer to the paper appeared in IEEE Transactions on Mobile Computing 12(2) for the complete set of results.

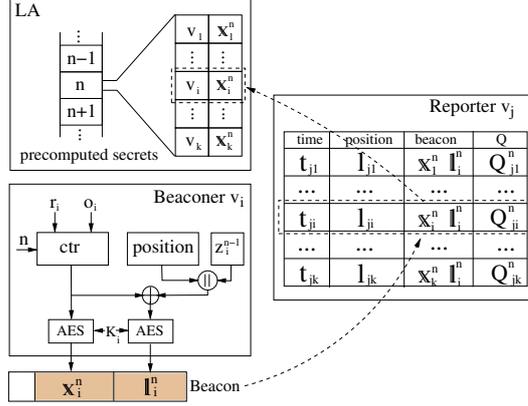


Figure 4.4: A-VIP procedures by beaconer, reporter and LA.

- in presence of unverified location claims, it grants the authority the capability to infer the actual position of malfunctioning or misbehaving vehicles;
- it does so by safeguarding drivers' privacy with respect to other vehicles participating in the network, and without any requirement for uninterrupted radio coverage from roadside infrastructure.

To achieve such goals, A-VIP leverages anonymous position beacons from vehicles, which prevent overhearing nodes from identifying or tracking their source, but still allow authorized third parties – sharing secret information with the beaconing vehicles – to perform such operations. Then, an authenticated reciprocal beacon reporting scheme grants the LA the possibility to verify the locations claimed by vehicles and infer unverified positions by efficiently solving an optimization problem.

4.2.1 A-VIP framework operation

The procedures in the A-VIP protocol involve three actors: the beaconer vehicle v_i broadcasting its position information that has to be verified, the reporter vehicle v_j that overhears the beacon by v_i and reports it to the LA, and the LA itself, which is in charge of performing the verification. Interactions are outlined by the schematic in Fig. 4.4. Specifically, each vehicle registers itself with the LA obtaining a registration triplet (K_i, r_i, o_i) where K_i is a short-term 128-bit AES symmetric key, and r_i, o_i are random integers. Vehicles acts as beaconers by periodically broadcasting anonymous beacons, which do not include the vehicle identifier and use a fresh random MAC-layer address [2]. The n -th beacon by vehicle v_i is assembled using the registration triplet and is composed of two parts:

- a time-dependent secret $\mathbf{x}_i(t) = E_{K_i}\{r_i + no_i\} = \mathbf{x}_i^n$, computed by incrementing the random-offset counter ($r_i + no_i$) and encrypting it with K_i using AES in counter mode (AES-CTR) [7]: the value of $\mathbf{x}_i(t)$ is only known to the vehicle and the LA, and allows the latter to verify the freshness of a beacon transmission and the identity of its originator;
- the encrypted current location announced by the vehicle $l_i^n = E_{K_i}\{(l_i^n || z_i^{n-1}) \oplus (r_i + no_i)\}$, where the plaintext location l_i^n is concatenated with a one-bit flag

z_i^{n-1} and XOR'ed with the plaintext counter value ($r_i + no_i$): both the flag and the counter allow to counter attacks to the A-VIP framework⁴.

Such a beacon structure allows for a fully anonymous information exchange, preventing overhearing and thus ensuring user privacy.

Vehicles also act as reporters for the beacons they receive. When a beacon issued by a vehicle v_i is correctly received by a vehicle v_j , the latter stores in a report table the following information: (i) the time t_{ji} at which the beacon is received; (ii) its own position l_{ji} at the time the beacon was received; (iii) the secret \mathbf{x}_i^n carried in the beacon; (iv) the encrypted position \mathbb{I}_i^n of v_i carried in the beacon; (v) a received signal quality field Q_{ji}^n . Reporters periodically upload their report table to the LA.

Upon reception of reports, the LA can:

- process them so as to estimate the probability $P_{i,s}^{(\mathcal{Z})}$ that a beaconer v_i is at position s , based on the data provided by a set \mathcal{Z} of reporters;
- compute such probabilities for all combinations of reporters, and leverage them to formulate an optimization problem that maximizes the consistency of reports by assigning a trustworthiness measure to each vehicle participating in the network;
- determine the final vehicle positions, either believing to those announced by trustworthy vehicles or estimating those of untrustworthy ones.

These three steps are detailed next.

Location estimation based on reports

By leveraging the information in the report tables, the LA can compute the probability that each beaoning vehicle was at a given location at a specific time instant. To that end, for each report table entry, the LA uses t_{ji} and \mathbf{x}_i^n to identify the vehicle v_i that sent the beacon reported in the entry, and thus to recover the matching symmetric key K_i . Using K_i , it decrypts the \mathbb{I}_i^n field, and extracts the announced position l_i^n and the flag z_i^{n-1} . If z_i^{n-1} is unset⁵ the LA stores the time t_{ji} , signal quality Q_{ji}^n , and locations of both the beaconer, l_i^n , and the reporter, l_j^n . Such information can be discretized (so that l_i^n and l_j^n map to geographical tiles s and t , while t_{ji} maps to a time step – dropped for the sake of clarity in the discussion below) and exploited to estimate the position of beaconer v_i at a given time instant. To do so, the LA needs a model of the propagation conditions⁶, which is a function $h(s, t, Q_{ji}) : \mathcal{S}^2 \times \mathbb{R} \rightarrow [0, 1]$ that, for any pair of geographical tiles (s, t) and any signal quality value Q_{ji} , provides the probability $\mathbb{P}(R_t^{(j)} | B_s^{(i)}, Q_{ji})$ that a beacon sent by v_i from tile s can be received by v_j in tile t , with the quality level Q_{ji} reported by v_j . By applying Bayes' theorem, the LA can use such values to compute the

⁴Details on the precise function of these values are available in the paper to appear in IEEE Transactions on Mobile Computing.

⁵If the flag z_i^{n-1} is set, the LA discards the entry as it is not reliable. Again, we refer the reader to the paper to appear in IEEE Transactions on Mobile Computing for details.

⁶The propagation model can be based on deterministic (e.g., ray-tracing), stochastic, or measurement-driven approaches, without changing the Q -aware procedure.

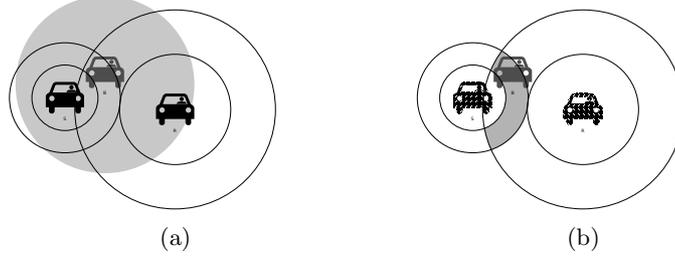


Figure 4.5: A-VIP position estimation example. (a) The shaded area represents the transmission range of v_i , while the annuluses denote the set of locations from which a beacon could be received by, respectively, v_k and v_j with the quality level indicated in their report. (b) the intersection of the annuluses represents the possible positions of v_i according to the A-VIP position estimation process on reports by v_k and v_j .

probability $\mathbb{P}(B_s^{(i)} | R_t^{(j)}, Q_{ji})$ that the beaconer was in tile s , given that the beacon was heard by v_j in tile t , with a quality level Q_{ji} . Specifically,

$$\mathbb{P}(B_s^{(i)} | R_t^{(j)}, Q_{ji}) = \frac{\mathbb{P}(R_t^{(j)} | B_s^{(i)}, Q_{ji}) \cdot \mathbb{P}(B_s^{(i)})}{\sum_{u \in \mathcal{S}} \mathbb{P}(R_t^{(j)} | B_u^{(i)}, Q_{ji}) \cdot \mathbb{P}(B_u^{(i)})}, \quad (4.1)$$

where $\mathbb{P}(B_x^{(i)})$, $x = s, u$, is the probability that the broadcasting vehicle v_i is in tile x at the time of transmission⁷.

Clearly, the same beacon, characterized by a single \mathbf{x}_i^t , may be reported by multiple vehicles v_j traveling in the proximity of v_i when the latter broadcasted it. Therefore, by repeating the procedure above on different report tables, the LA can obtain multiple expressions of $\mathbb{P}(B_s^{(i)} | R_t^{(j)}, Q_{ji})$, for different reporters v_j . The probabilities associated to any group of reporters $v_j \in \mathcal{Z}$ can be combined so as to improve the estimate that a beaconer v_i is at tile s at the considered time, as follows:

$$P_{i,s}^{(\mathcal{Z})} = \frac{\prod_{j: v_j \in \mathcal{Z}} \mathbb{P}(B_s^{(i)} | R_t^{(j)}, Q_{ji})}{\sum_{u \in \mathcal{S}} \prod_{j: v_j \in \mathcal{Z}} \mathbb{P}(B_s^{(i)} | R_t^{(j)}, Q_{ji})} \quad \forall s \in \mathcal{S}, \quad (4.2)$$

where \mathcal{S} is the set of all possible tiles. A simple example outlying the rationale for (4.2) is portrayed in Fig. 4.5, where two reporters, v_k and v_j , include different quality levels for a beacon received from v_i , whose communication range is highlighted in

⁷This value may depend on, e.g., road traffic information. In our performance evaluation, we will assume however a worst-case scenario where no such knowledge is available, and the probability is equally spread among all spatial locations.

Fig. 4.5(a). For simplicity, in the figure we considered that the area corresponding to the received quality value indicated by a reporter maps onto an annulus comprised in its reception range. Then, the set of possible locations of the beaconer is given by the intersection of the two annuluses, i.e., the shaded area in Fig. 4.5(b).

Trustworthiness level of beaconer vehicles

The location estimation technique implemented by (4.2) is leveraged to formulate an optimization problem, whose solution determines the level of trustworthiness $\gamma_i \in [0, 1]$ of each vehicle v_i . To that end, we first define a location probability measure $\Phi_{i,s}^{(\mathcal{R}_i)}$ that a vehicle v_i is at any tile $s \in \mathcal{S}$, which accounts for the (unknown) trustworthiness of all vehicles that reported the beacon sent by v_i , as follows:

$$\Phi_{i,s}^{(\mathcal{R}_i)} = \sum_{\mathcal{Z} \in \wp(\mathcal{R}_i)} \left(P_{i,s}^{(\mathcal{Z})} \prod_{j:v_j \in \mathcal{Z}} \gamma_j \prod_{k:v_k \in \mathcal{R}_i \setminus \mathcal{Z}} (1 - \gamma_k) \right), \quad (4.3)$$

where \mathcal{R}_i is the set of all reporters of the beacon sent by v_i at the considered time instant, and $\wp(\mathcal{R}_i)$ is the power set of \mathcal{R}_i , i.e., all possible subsets (proper and not) of reporters in \mathcal{R}_i . In words, the expression in (4.3) states that, if only the reporters in the subset \mathcal{Z} are trustworthy, which happens with probability $\prod_{j:v_j \in \mathcal{Z}} \gamma_j \cdot \prod_{k:v_k \in \mathcal{R}_i \setminus \mathcal{Z}} (1 - \gamma_k)$, then the probability that the beaconer v_i was in s is obtained by considering the reports sent by such vehicles ($v_j \in \mathcal{Z}$) and neglecting the others ($v_k \in \mathcal{R}_i \setminus \mathcal{Z}$). Note that, if $\gamma_j = 1 \forall j : v_j \in \mathcal{R}_i$, i.e., all reporters were to be trusted, the expression in (4.3) would reduce to $P_{i,s}^{(\mathcal{R}_i)}$, i.e., to the probability associated to the intersection of all the reported beacon receptions.

The optimization problem is then formulated as:

$$\max_{i:v_i \in \mathcal{V}} \sum \left(\gamma_i \Phi_{i,l_i}^{(\mathcal{R}_i)} + (1 - \gamma_i) \sum_{s \in \mathcal{S}} \Phi_{i,s}^{(\mathcal{R}_i)} \right). \quad (4.4)$$

The problem – whose variables are the γ_i 's – aims at maximizing the consistency among reports uploaded by vehicles in the network. Specifically, it considers each vehicle v_i in the network, whose set is denoted as \mathcal{V} , and sums up two terms. The first term corresponds to the case where v_i is correct (which happens with probability γ_i): in this case we consider the probability that v_i was in the (discretized) location l_i that it announced in its beacon, according to the reporters. The second term, instead, corresponds to the case where v_i cannot be trusted (which happens with probability $1 - \gamma_i$): in this case, we account for the probability that v_i could have been in any of the possible tiles, $s \in \mathcal{S}$. By simple variable transformation, the objective function in (4.4) can be expressed in *posynomial* form⁸. Posynomial problems can be reduced to a convex form and thus maximized in polynomial time [8].

The solution of (4.4) yields the level of trustworthiness γ_i for all vehicles $v_i \in \mathcal{V}$ that participate in the network.

⁸A posynomial is a function of the form $f(x_1, x_2, \dots, x_n) = \sum_{k=1}^K c_k x_1^{a_{1k}} \dots x_n^{a_{nk}}$ where all x_i and coefficients c_k are positive real numbers, and the exponents a_{ik} are real.

Final positions

Algorithm 1 Filling the set of trusted vehicles.

Require: $\gamma_i, \forall v_i \in \mathcal{V}$

- 1: $\mathcal{T} \leftarrow \emptyset$
 - 2: $\mathcal{T}' \leftarrow \emptyset$
 - 3: **repeat**
 - 4: $\mathcal{T} \leftarrow \mathcal{T}'$
 - 5: $v_i \leftarrow \arg \max_{h: v_h \in \mathcal{V} \setminus \mathcal{T}} \gamma_h$
 - 6: $\mathcal{T}' \leftarrow \mathcal{T} \cup \{v_i\}$
 - 7: **until** $(\exists v_k: v_i \in \mathcal{R}_k \wedge \max_{\mathcal{S}} P_{k,s}^{(\mathcal{T}'_k)} = 0) \vee \mathcal{T}' = \mathcal{V}$
 - 8: **return** \mathcal{T}
-

In order to determine the final positions of vehicles, the LA first needs to tell apart vehicles that can be ultimately trusted from those that cannot. There, a simple threshold mechanism on γ_i is not enough, since adversarial vehicles may easily influence in a negative way the γ_i of correct vehicles nearby. Therefore, the LA runs a dedicated algorithm, outlined in Alg. 1, whose goal is to determine the set $\mathcal{T} \subseteq \mathcal{V}$ of correct vehicles. At the outset, the LA initializes \mathcal{T} to the empty set (line 1). Then, at each step, it selects the vehicle v_i , in \mathcal{V} but not yet in \mathcal{T} , for which the probability to be trustworthy γ_i is the highest. It adds the vehicle to the set \mathcal{T}' , which is thus given by $\mathcal{T} \cup \{v_i\}$ (lines 5–6). If the information provided by v_i is consistent with the one provided by vehicles already in \mathcal{T} , then v_i is trusted as well and included in \mathcal{T} . Specifically, consistency is evaluated as follows: let us denote by \mathcal{T}'_k the set of vehicles that have reported the beacon sent by v_k and are in set \mathcal{T}' ; then v_i is consistent with nodes already in \mathcal{T} if adding its reports does not reduce to zero the probability $P_{k,s}^{(\mathcal{T}'_k)}$ for all vehicles and on all tiles v_k (line 7). If it is not trusted, v_i and all the reporters with a value of trustworthiness probability lower than γ_i are tagged as non-trusted, and the procedure ends.

Upon end of Alg. 1, the LA assigns to a position $\mathcal{L}_i = \{l_i\}$ to all $v_i \in \mathcal{T}$, i.e., it accepts the positions declared by trusted vehicles. Otherwise, if v_i is not in \mathcal{T} , the LA associates to it the set of possible locations $\mathcal{L}_i = \{s | P_{i,s}^{(\mathcal{R}_i \cap \mathcal{T})} > 0\}$, estimated according to the information provided by trusted reporters.

4.2.2 A-VIP framework performance evaluation

We evaluated the performance of A-VIP both in a synthetic scenario and in real-world live testbeds, representative of typical urban and suburban regions. A sample of the results obtained is provided in Fig. 4.6, where adversarial vehicles run false location attacks by advertising old locations rather than their current one, along with consistent cryptographic material.

Fig. 4.6(a) portrays the CDF of the trustworthiness probability γ for correct and adversarial vehicles, when the beaoning interval, named τ_b hereafter, is set to 1 s or 10 s. These simulation results consider attackers that include in their beacons 10 s-old positions, and show that A-VIP reliably separates two classes of nodes, assigning

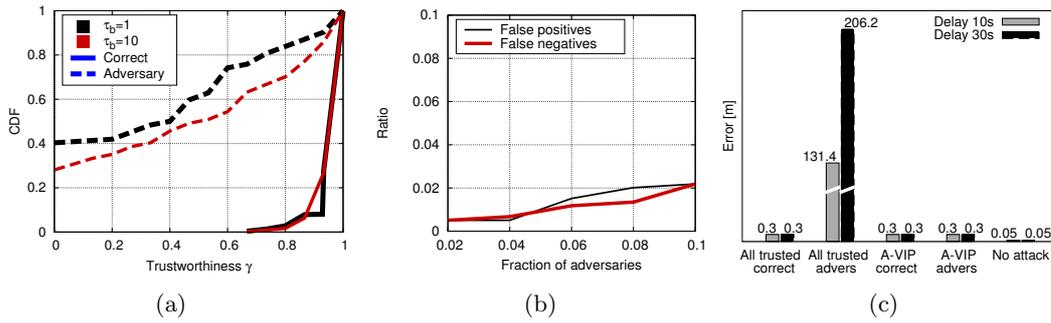


Figure 4.6: A-VIP performance evaluation, in presence of false location attacks. (a) Simulation: distribution of the trustworthiness probability γ , with 10% of adversaries announcing 10 s-old positions and a beaconing interval τ_b equal to 1 or 10 s. (b) Simulation: false positives and negatives versus fraction of attackers, with 10% of adversaries announcing 10 s-old positions and a beaconing interval τ_b of 10 s. (c) Testbed: error of the final positions determined by the LA, with 2 attackers out of 6 vehicles announcing 10 s- or 30 s-old positions and a beaconing interval τ_b of 1 s.

high γ values to correct vehicles and low γ values to adversarial ones. Only 5% to 10% of the attackers are assigned a high trustworthiness, and this mainly occurs for adversarial nodes that did not move significantly during the 10-second delay of the attack. In fact, the proper classification of correct and adversarial behaviors leads to extremely low false positives (i.e., attackers tagged as trustworthy) and false negatives (i.e., correct nodes tagged as adversarial), in Fig. 4.6(b). Both are below 2% of the vehicles, and the result holds for any fraction of adversaries between 2% and 10% of the total road traffic.

Fig. 4.6(c) shows instead a representative experimental result obtained in an urban testbed. The plot depicts the average error made by the LA in determining the final positions of vehicles, when $\tau_b = 1$ s and the false location attackers announce 10 or 30 s-old locations. The first two pairs of bars refers to the case where no verification is run by the LA, and the positions announced by all vehicles are simply trusted. We can observe that, while correct vehicles (“All trusted correct”) announce positions that are close to their real ones and thus yield low errors, adversarial vehicles (“All trusted advers”) succeed in announcing locations that are very distant from their actual ones. When A-VIP is employed, the error of adversarial vehicles (“A-VIP advers”) is drastically reduced to the level of that of correct vehicles (“A-VIP correct”). Moreover, the effectiveness of A-VIP is confirmed under different delays used by the attackers, and the errors it achieves are in all cases only slightly higher than those – negligible – incurred in when no attackers are present (“No attack”).

4.3 Conclusions of the Chapter

In this Chapter, I presented techniques that address the problem of verifying position information announced by communication-enabled vehicles. The problem was tackled from two different points of view. First, I considered a fully distributed scenario,

where one vehicle must autonomously decide on the correctness of the location information it receives from neighboring cars. Second, I moved to the perspective of a Location Authority having to validate positioning data from all communication-enabled vehicles within some geographical boundaries. In both environments, lightweight solutions were devised that leverage reciprocal reporting in order to tell apart correct and faulty or adversarial vehicles. Simulation and experimental results confirm that the proposed solutions are effective in identifying nodes advertising false positions, while keeping the probability of false positives low. In the case of the centralized approach, the solution also proved to be capable of precisely inferring the actual locations of misbehaving users.

Bibliography

- [1] 1609.2. “*IEEE Trial-Use Standard for Wireless Access in Vehicular Environments – Security Services for Applications and Management Messages*,” U.S. Department of Transportation, January 2006.
- [2] P. Papadimitratos, L. Buttyan, T. Holczer, E. Schoch, J. Freudiger, M. Raya, Z. Ma, F. Kargl, A. Kung, J.-P. Hubaux, “*Secure Vehicular Communications: Design and Architecture*,” IEEE Communications Magazine, 46(11), November 2008.
- [3] S. Čapkun, J.-P. Hubaux, “*Secure Positioning in Wireless Networks*,” IEEE Journal on Selected Areas in Communications, 24(2), 2006.
- [4] Nanotron Technologies, nanoLOC RTX factsheet, http://www.nanotron.com/EN/pdf/Factsheet_nanoLOC-NA5TR1.pdf.
- [5] B. Wiedersheim, M. Zhendong, F. Kargl, P. Papadimitratos, “*Privacy in Inter-Vehicular Networks: Why Simple Pseudonym Change is not Enough*,” IEEE WONS, Kranjska Gora, Slovenia, January 2010.
- [6] E. Schoch, F. Kargl, “*On the Efficiency of Secure Beaconing in VANETs*,” ACM WiSec, Hoboken, NJ, USA, March 2010.
- [7] W. Diffie, M. Hellman, “*Privacy and Authentication: An Introduction to Cryptography*,” Proceedings of the IEEE, 67(3), March 1979.
- [8] S.P. Boyd, L. Vandenberghe, “*Convex Optimization*,” Cambridge University Press, 2004.

Chapter 5

Perspectives

Conclusions on the topics addressed in this document are summarized at the end of each Chapter, or more thoroughly presented in dedicated *Discussion* sub-Sections when required. Therefore, I will not repeat myself here, but I will just recall what I consider to be the most significant highlights among the results obtained. Once more, all the researchers I had the luck to collaborate with over the past six years deserve a significant share of the merit for these achievements.

- We generated what can be considered today as the state of the art in the context of vehicular mobility datasets for networking studies. The dataset describes over 700,000 fine-grained car trips occurring during a typical working day within the 400-km² conurbation of Cologne, Germany. Its wide adoption by the research community, with tens of research groups worldwide currently employing the dataset for their research, is a proof of the relevance and timeliness of the work.
- We unveiled the poor availability, reliability and navigability of pure vehicular ad hoc networks. These conclusions were obtained by characterizing the protocol-independent topological properties of the vehicular network, and were derived in presence of realistic mobility. Our results clearly outline the limits of the VANET paradigm, and put into question a significant portion of the vast literature on the topic. Also, they pinpoint the need for some fixed infrastructure to support direct vehicular communication in most large-scale service use cases.
- However, in the reverse perspective, we showed that vehicle-to-vehicle communication can be very efficient in complementing a fixed infrastructure. In the optimal case, i.e., assuming perfect knowledge of the mobility of vehicles, peaks of 80% of the data traffic can be offloaded from the access network to the DSRC-based vehicular network. This result holds for very different data traffic models, namely the download of large-sized contents and the upload of small-sized floating car data.
- We devised simple and distributed solutions for data management under both traffic models above. We showed how easy-to-implement schemes can approx-

imate well the maximum performance bounds of the system. The encouraging results on optimal and practical data management in vehicular environments stimulated the community to undertake significant efforts on the topic. Such an impact of our work is also proven by the hundreds of papers that have been citing our research on this subject.

- We outlined the challenges and demonstrated the feasibility of secure positioning in vehicular environments, under both distributed and centralized approaches. Our studies helped drawing the attention of the research community on an overlooked yet critical aspect of vehicular networking.

Building on these highlights, I would also like to seize the opportunity offered by this concluding Chapter to hint at very recent activities and future perspectives of my research. These perspectives include what we can consider as medium-term research interests, which are presented in Sec. 5.1. There, methodologies and goals are quite well identified, and I expect to obtain conclusive and comprehensive results within the next few years on these topics. However, the same perspectives also allow sketching my long-term research plan, which is discussed in Sec. 5.2. In that case, the objectives are farther in time and thus definitely broader and less precise: the idea is thus to outline the general directions which I would like to steer my activities to, during the course of my career.

5.1 Medium-term interests

In the next few years, I would like to capitalize on the expertise I grew with my past works, but also to explore the possibilities offered by new emerging research fields.

On the one hand, I will keep on conducting research on topics related to the vehicular networking field. The fact that we are finally moving towards a mass commercialization of communication technologies in the automotive industry is a significant stimulus to stick with this subject. Indeed, this is an almost obliged feeling in front of projections that by 2016 automotive telematics is likely to be the largest contributor of the whole machine-to-machine market, with a 15.5 billion dollar value, that by 2025 90% of new cars will come with factory-fitted telematics [1], or that vehicle-to-vehicle communication is expected to avoid as much as 76% of potential multi-vehicle collisions involving cars [2].

On the other hand, I am growing an interest towards the analysis of mobile data traffic and its implications on the performance evaluation and design of protocols and architectures of mobile networks. My interest for a research field I have small experience of is two-fold. First, there exists today an unprecedented opportunity to obtain large datasets of people mobility and habits, e.g., from open data initiatives by municipalities [3, 4, 5, 6, 7, 8, 9], challenges [10, 11], social network applications such as Twitter or FourSquare, or direct collaborations with mobile network operators. Second, solutions for mobile data analysis that are novel and effective enough may have immediate practical applications over large deployments – which is not always the case with upcoming vehicular networking technologies.

I expect that these two topics will drive most of my medium-term work. Therefore, the remainder of this Section is dedicated to more detailed discussions on the specific research challenges and planned activities in the fields of vehicular networking (Sec. 5.1.1) and mobile data analysis (Sec. 5.1.2).

5.1.1 Vehicular networking

The definition of a reference set of mobile traffic traces for the evaluation of network solutions remains, in my opinion, one of the biggest near-term challenges concerning vehicular networking. Such a reference set should include mobility traces that are publicly available, realistic, highly detailed – from both space and time perspectives – and heterogeneous, i.e., they should capture varied geographical areas, traffic conditions and road network scales.

Despite all current efforts to gather real-world data [12, 13, 14], to develop dedicated simulation tools [15, 16, 17, 18], and to generate synthetic traces [19, 20, 21], we are still far from the objective. The lack of a reference set that is widely recognized and commonly adopted leads to performance evaluations of vehicular networks that are often unreliable and non-reproducible [22].

Therefore, an interesting medium-term objective is that of developing a reference set of road traffic traces that is specifically intended for the study of vehicular network protocols and architectures. In order to attain such a goal, one needs to (i) collect, (ii) classify, and (iii) provide the traces that will compose the reference dataset.

These activities, detailed next, are expected to be carried out in the context of national and international research projects. In this regard, early support will be provided by the European Union through the project FP7-PEOPLE-CIG n.630211 *ReFleX*, i.e., “Re-thinking the fundamentals of vehicular networking with transportation theory and complex network science”, for which I am currently acting as Scientist in Charge.

Dataset collection

As far as the collection phase is concerned, a sensible approach is that of carrying out two activities in parallel:

- the survey and collection of open-access datasets of vehicular mobility that are currently available or that will become so in the future;
- the generation of additional datasets from raw source data, in the form of, e.g., O/D matrices, intersection turning probabilities, traffic counts, or inter-arrival measurements.

Both activities are part of my current research efforts. We started to carry out the first activity in [24], where a survey of existing datasets is provided. The second activity heavily relies on availability of the raw source data. To that end, collaborations with transportation authorities are fundamental. As an example, recent interactions with the Dirección General de Tráfico (i.e., the office for the traffic management) in Madrid granted access to high-detail measurements of inter-arrival



(a) Buried induction loop (b) Measurement point A (c) Measurement point B (d) Map locations

Figure 5.1: Induction loop (a). Close-by view of two measurement points (b,c). Geographical location of the two points in the conurbation of Madrid, Spain (d).

times and speeds on highways that surround the city. The data was retrieved from pairs of induction loops buried under the concrete layer, as in Fig. 5.1(a). Fig. 5.1(b) and Fig. 5.1(c) provide photographic views of two measurement points on highways M40 and A6, whose locations are shown in Fig. 5.1(d). In order to understand the difficulty of accessing to this kind of data, we only mention that usually, the local controllers that receive induction loop data are programmed to supply coarse-grained information, aggregated over 60-second intervals. Clearly, that makes the data hardly useful for networking purposes: the controller configuration had thus to be changed specifically for our study, so as to obtain real-world traffic counts with fine-grained time, speed and lane information on each single transiting vehicle. To the best of our knowledge, this is the only traffic count dataset recorded to date with such a level of detail.

Source data such as that described above represent the foundations to the generation of reliable synthetic datasets of vehicular mobility. Indeed, the latter are obtained by feeding the real-world information to dependable simulators of road traffic, which are becoming increasingly available and an example of which is the SUMO tool employed in Chapter 2.

Dataset classification

Collecting or generating the datasets is not sufficient – they also need to be classified in a coherent and comprehensive manner. As a matter of fact, not all datasets are useful for evaluation of all vehicular networking solutions, and it is critical that a researcher can easily identify the single trace or the subset of traces that best fit a target use case. Apart from trivial distinctions (e.g., urban versus highway, city-scale versus district-scale, rush hour versus off-peak), a clear categorization requires understanding of the fundamental properties of each trace. To that end, I believe that the correct tool to employ is that of the characterization of the topological features of the vehicular network. As shown in the previous Chapters, such a characterization allows to understand the capabilities one can expect from the vehicular network, through connectivity measures that: (i) abstract the underlying mobility in a way that is easily intelligible from a networking standpoint; (ii) are independent from the overlaying network solution to be studied.

A significant example is provided in Chapter 2, yet the same process can be systematically repeated for any road traffic dataset. For instance, in the case of multiple traces (collected over two highways at 8.30 am and 11.30 am, on four different days)

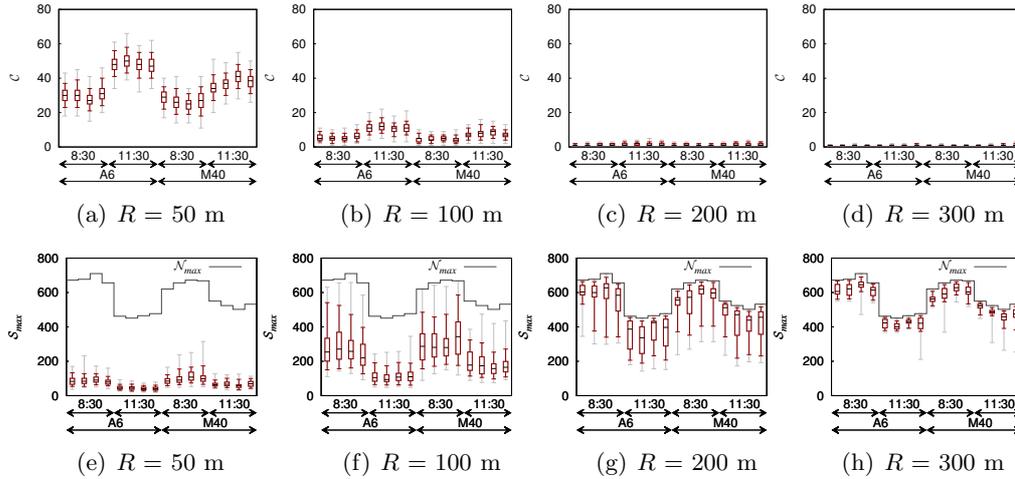


Figure 5.2: Distributions of the number of components \mathcal{C} (top), and of the largest component size \mathcal{S}_{max} (bottom), for each trace under different values of R .

generated from the high-detail traffic counts collected in Madrid and mentioned above, we obtain the results shown in Fig. 5.2. The plots refer to network-level analysis and show the distributions of the number of components \mathcal{C} (top plots) and of the size of the largest component \mathcal{S}_{max} (bottom plots), for various values of the communication range R (on different columns). In every plot, each candlestick refers to one mobility trace: the lowest and highest values (in grey) are the minimum and maximum values that the metric attains; the inner error bars (in red) depict the first and ninth deciles; the box highlights the first and third quartiles (in red) around the median value (in black). Also, the step function in the \mathcal{S}_{max} plots is the maximum value \mathcal{N}_{max} of \mathcal{N} , observed in all instances of the respective trace. It thus represents the upper bound to \mathcal{S}_{max} : the closer is \mathcal{S}_{max} to \mathcal{N}_{max} , the closer the vehicular network is to a fully connected single component.

These plots make it immediately clear which is the level of network connectivity, or fragmentation, in each dataset. E.g., for $R = 50$ m, there exist on average 30-50 disconnected components, the largest of which only comprises around one tenth of the vehicles. As R grows, however, the network fragmentation is reduced, and more nodes join the largest component. When $R = 300$ m almost all vehicles belong all the time to one single component, as in Fig. 5.2(d) and Fig. 5.2(h). Also, we observe that the connectivity metrics are very similar on different days of the week, which implies that we can expect the network to have consistent topological properties on Monday to Friday. Instead, the network appears more fragmented on A6 than on M40, and, in both cases, the connectivity is worse at 11:30 a.m. than at 8:30 a.m.

Overall, the proposal is that of classifying road traffic datasets according to the topological features they induce on the vehicular network. That will allow to neatly identify, e.g., traces where vehicle-to-vehicles connectivity is pervasive, spotty or quasi-absent, where very large or small clusters of vehicles are present, where vehicles find themselves into neighborhoods of hundreds, tens or just a few other cars, etc. In fact, in the context of networking studies, these metrics are much more representative than popular scenario features such as, e.g., the number of vehicles

present in the trace or the layout of the road network.

Dataset provisioning

The final step is that of providing easy and public access to the vehicular mobility traces that have been collected and classified according to their topological features. To that end, the aim is that of developing a single repository, which will become the reference point to retrieve datasets of interest. That is similar to what happens, e.g., for wireless network datasets with the successful initiative of CRAWDAD [23].

5.1.2 Mobile data analysis

As far as mobile data analysis is concerned, I am mainly interested in developing complex network science techniques that allow the automated mining of data generated by mobile users and collected by network, platform or service providers. I believe that such techniques may reveal paramount in letting the mobile network architecture accommodate the always growing capacity demand in densely populated metropolitan areas. As such they may end up complementing other envisioned solutions, based on, e.g., cognitive radio [25], femtocell and WiFi offloading [26], cooperative relaying [27], or re-design of the access network architecture entirely [28].

To that end, I intend to explore two complementary approaches. The first takes the operator's perspective: mobile data is leveraged to understand where, when, and how the access network resources are consumed by mobile users. The second considers an individual user's viewpoint: mobile data is leveraged to understand how each user behaves in terms of mobility patterns, interaction with other users, and mobile service consumption. These two perspectives are discussed in detail in the remainder of this Section.

Also in this case, the research activities are expected to be carried out in the framework of national and international collaborative projects. Currently, research on the operator's perspective mentioned above is supported by the National Research Agency in France, through the project ANR-13-INFR-0005 *ABCD*, i.e., "Adaptive Behavior and Cloud Distribution", where I am acting as work package leader. Research on the mobile user's perspective is instead supported by the European Union FP7 ERA-NET program, through project CHIST-ERA-2012 *MACACO*, i.e., "Mobile context-Adaptive Caching for COntent-centric networking", where I am the Principal Investigator for the CNR partner.

Mobile network usage profiling

Individual mobile users consume network resources in significantly different ways, depending on the time at which they access the network and on the location where they do so. When aggregating user behaviors, very diverse macroscopic network utilization profiles emerge that vary over space and time. Therefore, traditional static capacity planning is hardly capable of coping with such a diversity and results in over-dimensioning and under-utilization of resources. On the contrary, accurately understanding the spatiotemporal dynamics of mobile customers' demand allows

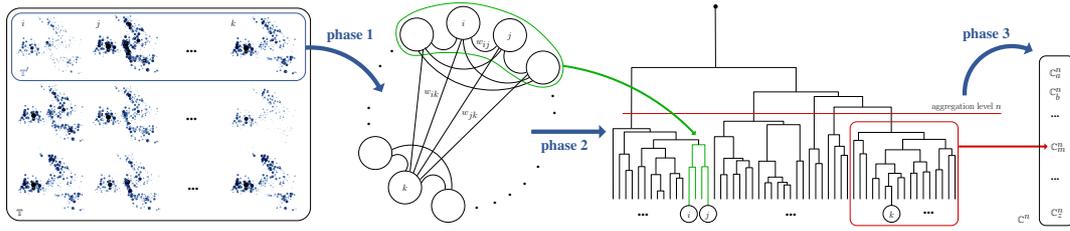


Figure 5.3: Workflow of the framework for the definition of categories of network usage profiles. Phase 1: construction of the snapshot graph from snapshots (portrayed here as geographical plots of the mobile traffic volume) in the training set. Phase 2: iterative aggregation of graph vertices into a dendrogram structure. Phase 3: identification of the clustering level n granting the maximum separation between the groups of snapshots. The resulting clusters are mapped to network usage profile categories.

operators to allocate resources more efficiently. This is especially important in urban areas, where vast user displacements occur over short space and time scales, forcing the mobile network capacity to rapidly adapt to – or even to anticipate – demand fluctuations.

We started addressing this topic by focusing on the the problem of classifying mobile demand profiles in large-scale mobile datasets provided by a network operator. We propose an automated, parameter-free framework that allows constructing categories of call profiles from a training mobile dataset and classifying network usages accordingly. As a by-product of such operations, the framework can tell apart typical and outlying behaviors in the network.

Specifically, the framework runs on *snapshots* of the mobile demand extracted from raw mobile data. As the name suggests, a snapshot is a representation of the load generated by mobile users on the access network at a given time instant. Apart that, we do not impose any constraint on the way snapshots are defined: they can describe the traffic volume at every second or averaged over longer time intervals, at each base station or aggregated over larger geographical areas, and for one or multiple types of services (e.g., voice calls, short text messages, Internet-based applications, etc.). Snapshots are processed by the framework in four phases. The first three phases aim at defining a limited number of network usage categories by analyzing a training set of snapshots, and their workflow is depicted in Fig. 5.3. The first phase builds a clique graph where each snapshot is mapped to a node, and each edge is weighted according to the similarity between the snapshots it links. The second phase leverages the graph to progressively aggregate snapshots and generate a dendrogram structure. In the third phase network usage categories are outlined by applying to the dendrogram stopping rules that determine which aggregation stage yields the most neat classification of profiles. Finally, in the fourth phase, not shown in Fig. 5.3 for the sake of clarity, we consider the obtained categories as the base structure for the classification of additional snapshots.

Evaluation on a 5-month mobile dataset with 300 million entries collected in Abidjan, Ivory Coast, showed that our framework can successfully identify call profile

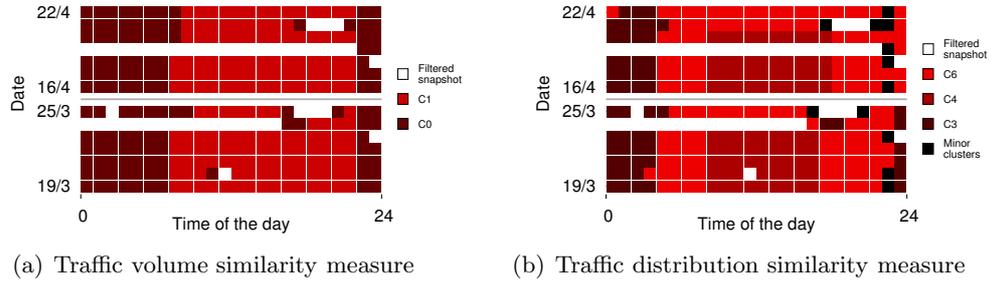


Figure 5.4: Classification of two weeks of snapshots into mobile traffic profile categories, according to two different similarity measures. Each square represents a 1-hour snapshot, whose category maps to a color. Each row maps to one day. Empty squares are hours for which data was not available.

categories that yield meaningful social properties. Also, it is capable of correctly classifying network usages within the aforementioned categories, and of detecting unusual behaviors in the mobile demand, which we show to have clear social origins.

Samples of the structure of the categories identified using two different snapshot similarity measures in shown in Fig. 5.4. We note that the two categories identified by the first measure clearly separate times with a lower activity, i.e., hours between 22:00 and 7:00, and times with a higher traffic, i.e. hours between 8:00 and 21:00. More interestingly, the second metric identifies three major clusters. The first category includes the snapshots of the night hours, between 23:00 and 4:00, characterized by a low traffic generated in the residential areas of the city. The second category includes daytime snapshots from the weekdays, i.e., hours between 10:00 and 17:00, Monday to Friday: these snapshots show higher mobile traffic in the office and university areas. The third major category contains snapshots from weekend days, i.e., Saturdays and Sundays, as well as from early morning (5:00–9:00) and evening (19:00–22:00) hours of weekdays: the corresponding network usage is that of a high traffic volume generated in the residential areas. Minor clusters also appear, which include a very small number of snapshots each.

Overall, even if they are just the outcome of early-stage studies, these results already hint at the significant potential applications that automated data mining frameworks can have for network planning, dimensioning and management.

Mobile user behavior profiling

Network-centric analyses do not represent the only possibility. Also the perspective of the mobile user can be of interest towards the objective of finding new ways to manage the increased data usage and to improve the level of service required by the new wave of smartphones applications.

In this case, the focus is on data offloading mechanisms that take advantage of context and content information. The intuition is that if it is possible to extract and forecast the behavior of individual mobile network users in the three-dimensional space of time, location and interest (i.e. “when”, “where” and “what” users are pulling data from the network), it is possible to derive efficient data offloading pro-

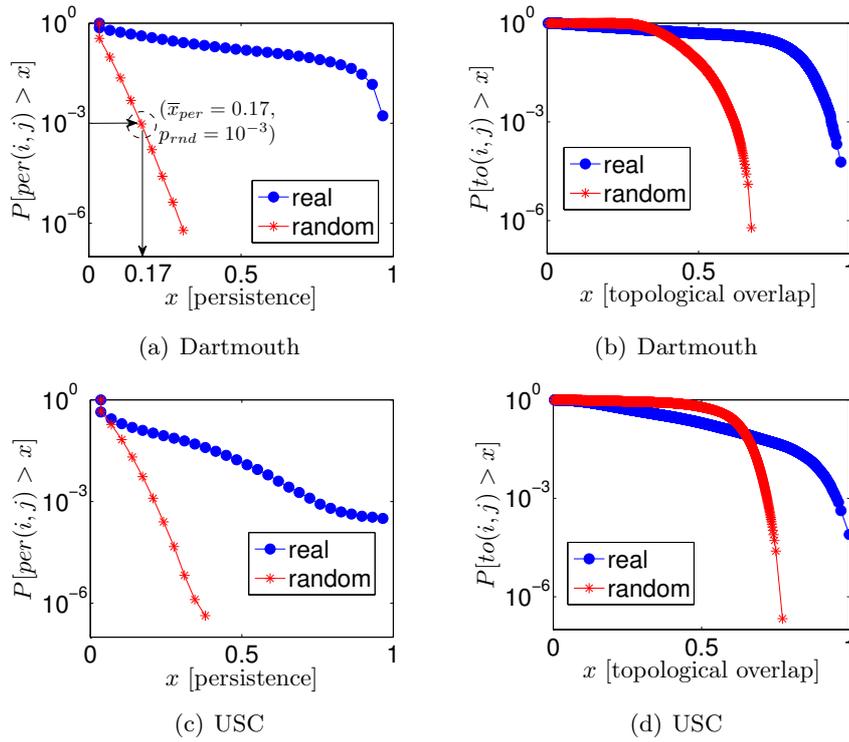


Figure 5.5: The complementary cumulative distribution function (CCDF) of the edge persistence (a,c) and topological overlap (b,d) for the real-world contact graph and its random equivalent. Results refer to two different four-week datasets, collected at Dartmouth (a,b) and USC (c,d) campuses.

ocols. Such protocols would pre-fetch the identified data and cache it at the network edge at an earlier time, preferably when the mobile network is less charged or offers better quality of service. Caching can be done directly at the mobile terminals, as well at the edge nodes of the network (e.g., femtocells or wireless access points).

Attaining the goal above means bringing together social wireless networking, opportunistic communications and content networking, so as to address several issues:

- the characterization of the mobility and social relationships of individual users, so as to model and anticipate their movement patterns and pairwise contacts;
- the derivation of appropriate models of the correlation of user interests with their mobility and social relationships.
- the design of efficient data offloading algorithms that build on the models developed at the two previous points. These include pre-fetching mechanisms that both improve the perceived quality of service of the mobile user and noticeably reduce peak bandwidth demands at the cellular network.

Initial efforts to address the first issue above are on-going. In particular, we started developing a technique to analyse mobile data and tell apart contacts among

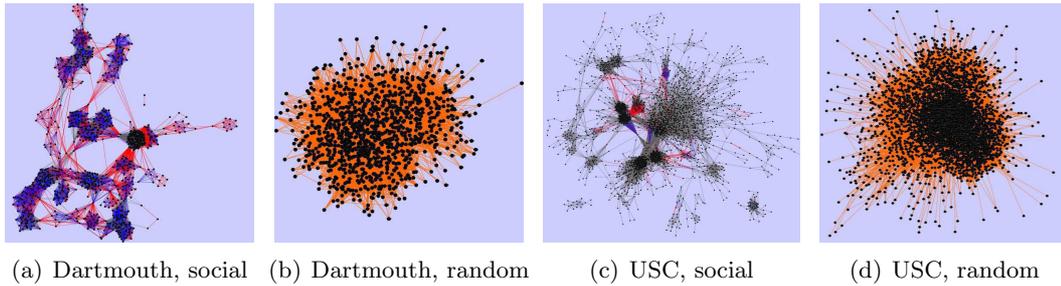


Figure 5.6: Snapshots of the Dartmouth (a,b) and USC (c,d) contact graphs, considering only the social (a,c) edges or random (b,d) edges.

users that are “social” (i.e., driven by routinary habit) and “random” (i.e., driven by pure fatality). The proposed technique extracts from the mobile data a temporal graph where nodes map to users, and edges to contacts. Then, it builds a random equivalent – a second contact graph that has the same node degree distribution of the first, but whose edges are drawn randomly. Finally, it compares the persistence and topological overlap of edges in the two graphs to identify those that are due to “social” interactions. As shown in Fig. 5.5, there is a negligible probability that the contacts in the random equivalent show persistence and topological overlap higher than a given threshold, which depends on the mobile dataset but can be easily identified from the associate probability distribution. By removing edges with persistence and topological overlap measures below that threshold, random contacts are discarded, and the actual social structure of the user network is unveiled. That is clear from Fig. 5.6, where we separate the contact graphs built on interactions classified as “social” and “random” in two sample mobile datasets.

Again, these results are preliminary. Still they show promise in improving our capability to understand mobile user movement and social interactions. In turn, a better knowledge of these features – along with a proper characterization of mobile service consumption – will be critical to devise effective solutions that can improve the way the access network accommodates mobile usages.

5.2 Long-term perspectives

On a longer term, my research objective is that of carrying out interdisciplinary research at the interface of multiple fields. I believe that unexplored interactions among different scientific disciplines yield a huge potential for innovation – especially when these disciplines are already independently addressing a same subject, from different perspectives and using diverse tools.

Clearly, my main focus will remain that of mobile telecommunication networks. However, I find attractive the idea of merging traditional approaches to mobile networking with those developed and employed by two other disciplines, i.e., big data science and complex network science.

On the one hand, I expect that the current trend in the growing availability of large amounts of (more or less open) data will not stop, rather it will accelerate in

the future. We may be soon submerged by vast quantities of useful information, hidden into orders-of-magnitude vaster useless data. A significant portion of such big data will concern the telecommunication field, in its many and varied facets: it will thus be critical to develop techniques that can dig into the raw data and extract information of interest to network, platform or service operators, such as precise profiles of network usages, mobile data traffic demands, and mobile user behaviors. Ideally, that result is to be achieved in a completely automated manner. To that end, the sensible choice is to establish strong interactions with computer scientists and statisticians and acquire expertise in managing and mining very large databases, i.e., in what is starting to be referred today as big data science.

On the other hand, even the portion of useful information extrapolated from the aforementioned big telecommunication data will be vast. Therefore, there is small hope that one will be able to understand such information by trivial approaches, e.g., simple visual inspection. Instead, there is a need to structure the data, and analyse it in a sound, comprehensive and, again, automated fashion. In this regard, graph representations are very powerful structures that can efficiently abstract most types of information. And, complex network science is an emerging discipline in physics that focuses on the analysis of very large graphs, so as to characterize their topology and features. The original and diverse methodologies that are being developed within the complex network science domain have been already used to understand and model a number of real-world systems, and may prove essential also in the case of mobile networking field.

I believe that tools developed for the study of big data and complex networks will reveal paramount to understand the lacks, overprovisionings and general operation of future mobile communication systems. The latter, in turn, are the essential information needed for an effective design of mobile network algorithms, protocols and architectures. Ultimately, I hope to be able to bring together expertises on these different disciplines into a research group focused on large-scale mobile networking. Ideally, the group will develop strong, fruitful collaborations with partners in academia and in the telecommunication industry. In that sense, INSA Lyon and Inria will be preferred candidates for cooperations. Overall, I would like the research carried out by the group to have an impact on network solutions deployed in the next generations of access and autonomous mobile telecommunication systems.

Bibliography

- [1] Telematics Update, “*The Connected Car in Europe: Gaining market share with telematics*,” White Paper, November 2013.
- [2] US Government Accountability Office, “*Vehicle-to-Vehicle Technologies Expected to Offer Safety Benefits, but a Variety of Deployment Challenges Exist*,” Report to Congressional Requesters, November 2013.
- [3] OpenBaltimore, <https://data.baltimorecity.gov>.
- [4] Dublinked, <http://www.dublinked.ie/datastore/datastore.php>.

- [5] Londo Datastore, <http://data.london.gov.uk>.
- [6] Roma Capitale Open Data, <http://dati.comune.roma.it>.
- [7] ParisData, <http://opendata.paris.fr>.
- [8] SmartData Grand Lyon, <http://smartdata.grandlyon.com>.
- [9] Rennes métropole en acces libre, <http://www.data.rennes-metropole.fr>.
- [10] Orange Data for Development Challenge, <http://www.d4d.orange.com>.
- [11] Telecom Italia Big Data Challenge, www.telecomitalia.com/tit/en/bigdatachallenge.html.
- [12] H.-Y. Huang, P.-E. Luo, M. Li, D. Li, X. Li, W. Shu, and M.-Y. Wu, “*Performance evaluation of swnet with real-time traffic data*,” IEEE VTC, 2007.
- [13] M. Doering, T. Pögel, W.-B. Pöttner, and L. Wolf, “*A new mobility trace for realistic large-scale simulation of bus-based DTNs*,” ACM CHANTS, 2010.
- [14] J. Yuan, Y. Zheng, X. Xie, and G. Sun, “*Driving with knowledge from the physical world*,” ACM KDD, 2011.
- [15] D. Krajzewicz, G. Hertkorn, C. Rössel, and P. Wagner, “*SUMO (Simulation of Urban MObility)*,” MESM, 2002.
- [16] J. Härri, M. Fiore, F. Filali, and C. Bonnet, “*Vehicular mobility simulation with VanetMobiSim*,” Simulation, 87(4):275–300, 2011.
- [17] C. Sommer, R. German, and F. Dressler, “*Bidirectionally coupled network and road traffic simulation for improved IVC analysis*,” IEEE Transactions on Mobile Computing, 10(1):3–15, 2011.
- [18] EU FP7 iTetris, <http://ict-itetris.eu>.
- [19] B. Raney, N. Cetin, A. Vollmy, M. Vrtic, K. Axhausen, and K. Nagel, “*An agent-based microsimulation model of swiss travel: first results*,” Networks and Spatial Economics, 3:23–41, 2003.
- [20] S. Uppoor, O. Trullols-Cruces, M. Fiore, and J. Barcelo-Ordinas, “*Generation and analysis of a large-scale urban vehicular mobility dataset*,” IEEE Transactions on Mobile Computing, to appear.
- [21] Y. Pigné, G. Danoy, and P. Bouvry, “*A vehicular mobility model based on real traffic counting data*,” Nets4Cars/Nets4Trains, 2011.
- [22] S. Joerer, C. Sommer, and F. Dressler, “*Toward reproducibility and comparability of IVC simulation studies: a literature survey*,” IEEE Communications Magazine, 50(10):82–88, 2012.
- [23] CRAWDAD: A Community Resource for Archiving Wireless Data At Dartmouth, <http://crawdad.cs.dartmouth.edu>.

- [24] S. Uppoor, M. Fiore, J. Härri, “*Synthetic mobility traces for vehicular networking*,” Vehicular Networks: Models and Algorithms, H. Labiod, A.-L. Beylot (Editors), Wiley, 2013.
- [25] B. Wang, R.K.J. Liu, “*Advances in cognitive radio networks: A survey*,” IEEE Journal of Selected Topics in Signal Processing, 5(1):5–23, 2011.
- [26] K. Lee, J. Lee, Y. Yi, I. Rhee, S. Chong, “*Mobile Data Offloading: How Much Can WiFi Deliver?*” IEEE/ACM Transactions on Networking, 21(2):536–550, 2013.
- [27] S. Kadloor, A. Raviraj, “*Relay selection and power allocation in cooperative cellular networks*,” IEEE Transactions on Wireless Communications, 9(5):1676–1685, 2010.
- [28] S. Liu, J. Wu, C.H. Koh, V.K.N. Lau, “*A 25 Gb/s(/km²) urban wireless network beyond IMT-advanced*,” IEEE Communication Magazine, 49(2):122–129, 2011.

Appendix A

Curriculum Vitae

MARCO FIORE



Main address

CNR – IEIIT
Corso Duca degli Abruzzi, 24
10129 Torino, Italy
(+39)0110905434
marco.fiore@ieit.cnr.it

Alternative address

CITI – Inria
20, avenue Albert Einstein
69621 Villeurbanne, France
(+33)472437319
marco.fiore@inria.fr

Academic background

<i>Mar</i> 2013	Researcher at CNR-IEIIT, Torino, Italy
<i>Jan</i> 2012	Co-founder of the Inria Urbanet team, Lyon, France
<i>Sep</i> 2009	Affiliation to Inria, within the SWING team, Lyon, France
<i>Sep</i> 2009	Tenured assistant professor at Institut National de Sciences Appliquées (INSA) de Lyon, France
<i>Sep</i> 2008	Visiting Postdoctoral fellow at Universitat Politècnica de Catalunya, Spain
<i>Apr</i> 2008	Postdoctoral fellow at Politecnico di Torino, Italy
<i>Apr</i> 2008	PhD degree in Electronics and Telecommunications Engineering from Politecnico di Torino, Italy
<i>Sep</i> 2006	Visiting researcher at Rice University, TX, USA
<i>Jan</i> 2005	PhD candidate in Electronics and Telecommunications Engineering at Politecnico di Torino, Italy
<i>Aug</i> 2004	Associate researcher at Politecnico di Torino, Italy
<i>Jul</i> 2004	Master of Science degree in Computer Science from Politecnico di Torino, Italy
<i>Dec</i> 2003	Master of Science degree in Computer Science from University of Illinois at Chicago, IL, USA

Teaching experience

2010–13	" <i>Network Performance Evaluation</i> ", INSA Lyon, France (in French), 250 hours
2010–12	" <i>Networks of the Future</i> ", Université de Lyon, France (in English), 24 hours
2010–12	" <i>Advanced Wireless Networks</i> ", INSA Lyon, France (in English), 40 hours
2009–12	" <i>Network Architectures, Services and Protocols</i> ", INSA Lyon, France (in French), 200 hours
2009–11	" <i>Introduction to Networking</i> ", INSA Lyon, France (in French), 90 hours
2009	" <i>Ad-hoc Wireless Networks</i> ", Politecnico di Torino, Italy (in English), 16 hours
2007–09	" <i>Evolution of Wireless Communication Networks</i> ", Politecnico di Torino, Italy (in Italian), 70 hours
2005–06	" <i>Introduction to Telecommunications Networks</i> ", Politecnico di Torino, Italy (in Italian), 50 hours

Research activities and interests

Vehicular networking

- Modeling of vehicular traffic in urban environments and analysis of the impact of road traffic dynamics on network protocols and architectures built on standards for vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I), including IEEE 802.11p, IEEE 1609.x, ETSI ITS G5, ISO CALM
- Management of content download, Floating Car Data (FCD) upload and information propagation from/by/among vehicular users, towards opportunistic relaying and cellular offloading via Direct Short-Range Communication (DRSC)
- Deployment of roadside units (RSUs) for vehicular network access.

Mobile networking

- Profiling, modeling and predicting user mobility and service usages in large-scale wireless access networks (GPRS, UMTS, WiFi, LTE), towards efficient resource allocation
- User position verification in mobile networks
- Inference of user relationships from their social contacts, towards efficient content dissemination in opportunistic Disruption-Tolerant Networks (DTN)
- Cooperative and distributed information retrieval, sharing and caching in mobile ad-hoc networks (MANETs)
- Experimental evaluation of client-driven techniques for the support of mobile users in urban Wireless Mesh Networks (WMNs), analysis and improvement of multihop routing in WMNs.

Advised students

PhD students

- Diala Naboulsi, PhD student, “*Human mobility - an urban networking perspective*”, September 2012 - October 2015 (planned)
- Sandesh Uppoor, PhD student, “*Understanding and exploiting mobility in wireless networks*”, October 2010 - December 2013 (now at Orange Labs, France)
- Pedro Olmo Stancioli, visiting PhD student, “*Telling apart social and random relationships in wireless networks*”, February 2011 - August 2011 (now at Universidade Federal de Minas Gerais).

Graduate students and interns

- 18 among graduate students and interns, mainly on four-month projects.

Professional activities

Chairing service

- Co-chair for IEEE/IFIP Wireless Days for the Track Wireless Models and Simulations, Valencia, Spain, November 2013
- Co-chair for IEEE Vehicular Technology Conference (VTC-Spring) for the Track Wireless Networks, Access Control, and Resource Management, Dresden, Germany, June 2013
- Co-chair for the First Workshop on Urban Networking (UrbaNe), co-located with the ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT), Nice, France, December 2012
- Co-chair for the Journées Spécifiques Pôle ResCom Réseaux Véhiculaires (ReVe), Lyon, France, November 2010.

Invited talks

- “Complex network science: the next big thing”, panel talk at IEEE NetSciCom, Toronto, Canada, May 2014
- “Large-scale urban vehicular networks: mobility and connectivity”, invited talk at the Hamilton Institute, Maynooth, Ireland, October 2012
- “Back to basics: road traffic vehicular networking”, keynote speech at the VANETs from Theory to Practice Workshop (VTP), S.Francisco, CA, USA, June 2012
- “Vehicular Networking: History, State of the Art and Interdisciplinary Aspects”, tutorial at the International Conference on Intelligent Transportation Systems and Telecommunications (ITST), S.Petersburg, Russia, August 2011, with Jérôme Härri
- “Vehicular mobility modeling”, tutorial at the RESCOM Summer School, Giens, France, June 2010, with Jérôme Härri.

PhD thesis committees

- Nadia Haddadou, Université Paris-Est, 2014 (examiner)
- Razvan Stanica, ENSEEIHT, 2011 (examiner)
- Valentina Martina, Politecnico di Torino, 2011 (reviewer)
- Saed Tarapiah, Politecnico di Torino, 2010 (reviewer)

Technical Program Committee service

- TPC member for the IEEE Global Communications Conference (GLOBECOM), Austin, TX, USA, December 2014
- TPC member for the Extreme Conference on Communication (ExtremeCom), Galápagos Islands, Ecuador, August 2014
- TPC member for the IEEE International Conference on Sensing, Communication and Networking (SECON), Singapore, June 2014
- TPC member for the IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM), Sydney, Australia, June 2014
- TPC member for the IEEE International Conference on Communications (ICC), Sydney, Australia, June 2014
- TPC member for the IEEE International Conference Wireless On-demand Network Systems and Services (WONS), Obergurgl, Austria, April 2014
- TPC member for the International Conference on Connected Vehicles and Expo (ICCVE), Las Vegas, NV, USA, December 2013
- TCP member for the IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), Lyon, France, October 2013
- TPC member for the IEEE Vehicular Technology Conference (VTC-Fall), Las Vegas, NV, USA, September 2013
- TPC member for the Extreme Conference on Communication (ExtremeCom), Iceland, August 2013
- TPC member for the IEEE International Conference on Computer Communication Networks (ICCCN), Nassau, Bahamas, July 2013
- TPC member for the IEEE International Conference on Sensing, Communication and Networking (SECON), New Orleans, LA, USA, June 2013
- TPC member for the IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM), Madrid, Spain, June 2013
- TPC member for the IEEE International Conference Wireless On-demand Network Systems and Services (WONS), Banff, Canada, March 2013
- TCP member for the IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), Barcelona, Spain, October 2012
- TPC member for the IEEE International Conference on Computer Communication Networks (ICCCN), München, Germany, July 2012

- TPC member for the ACM International Workshop on Hot Topics in Planet-scale Measurement (HotPlanet), Low Wood Bay, UK, June 2012
- TPC member for the International Workshop on Vehicular Communications and Applications (VCA), Ayia Napa, Cyprus, June 2012
- TPC member for the IEEE International Conference Wireless On-demand Network Systems and Services (WONS), Courmayeur, Italy, January 2012
- TPC member for the IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM), S.Francisco, CA, USA, June 2012
- TPC member for the Extreme Conference on Communication (ExtremeCom), Zurich, Switzerland, March 2012
- TPC member for the Extreme Conference on Communication (ExtremeCom), Manaus, Bresil, September 2011
- TPC member for the IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Toronto, Canada, September 2011
- TPC member for the IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM), Lucca, Italy, June 2011
- TPC member for the International Conference Wireless On-demand Network Systems and Services (WONS), Bardonecchia, Italy, January 2011
- TPC member for the IEEE Workshop on Intelligent Vehicular Networks (InVeNET), S.Francisco, CA, USA, November 2010
- TPC member for the International Symposium on Wireless Personal Multimedia Communications (WPMC), Brest, France, September 2011
- TPC member for the International Wireless Communications and Mobile Computing Conference (IWCMC), Caen, France, June 2010
- TPC member for the IEEE International Symposium on Wireless Vehicular Communications (WiVeC), Taipei, Taiwan, May 2010
- TPC member for the IEEE Workshop on Intelligent Vehicular Networks (InVeNET), Macau, China, October 2009
- TPC member for the International Symposium on Wireless Personal Multimedia Communications (WPMC), Sendai, Japan, September 2009
- Shadow TPC member for the ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT), Madrid, Spain, December 2008
- TPC member for the IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Cannes, France, September 2008
- TPC member for the International Symposium on Wireless Personal Multimedia Communications (WPMC), Lapland, Finland, September 2008.

Administrative responsibilities

- Elected member of the council of the Telecommunication Department at INSA Lyon, 2011-2013
- Faculty referent for the Projets de Fin d'Etudes (master thesis) defended at the Telecommunication Department at INSA Lyon, 2011-2013
- Member of the selection committee for a tenured assistant professor position at INSA Lyon, n. 27MCF540, April 2012.

Projects*Collaboration network*

I participated in several collaborative research projects, listed next. Within their framework, I got in touch with a large number of international partners, the most notable being: Politecnico di Torino, CNR, Fiat Research Center (Italy); Universitat Politècnica de Catalunya (Spain); INSA Lyon, Inria, Orange Labs, Alcatel-Lucent, Université Pierre et Marie Curie, ENSEEIHT, Eurecom (France); University of Birmingham (UK); SUPSI (Switzerland).

International projects

- Scientist in Charge in ReFleX – Re-thinking the fundamentals of vehicular networking with transportation theory and complex network science, FP7-PEOPLE CIG 2014-2018
- CNR–IEIIT partner leader in MACACO – Mobile context-Adaptive Caching for Content-centric networking, CHIST-ERA 2013-2016, Dr. A. Carneiro Viana, Inria, France
- Participation to EuroNF – Design and Engineering of the Next Generation Internet, EU NoE 2008-2010, JRAS.8 - Information Mobility in Ad Hoc Wireless Networks, Prof. J.M. Barcelo, UPC, Spain
- Participation to EuroNGI – Design and Engineering of the Next Generation Internet, EU NoE 2003-2006, JRA.2.1 - Mechanisms and protocols for controlled bandwidth sharing, Prof. M. Johansson, KTH, Sweden
- Participation to NEWCOM – Network of Excellence in Wireless Communications, EU NoE 2005-2006, WPRA - Ad Hoc and Sensor Networks, Prof. S. Palazzo, Università di Catania, Italy.

National projects

- Work Package leader in ABCD – Adaptive Behavior and Cloud Distribution, ANR INFRA 2013-2016, Dr. R. Langar, LIP6–UMPC, France
- Inria partner leader in SelfNet – Self-organizing Networks, Research action of the Inria / Alcatel-Lucent Bell Labs Common Laboratory 2010-2013, Dr. L. Rouillet, Alcatel-Lucent Bell Labs, France
- Participation to VICSUM – Vehicle-to-Vehicle-to-Infrastructure Communication for Sustainable Urban Mobility, Bando Regionale 2007-2008, Prof. C. Casetti, Politecnico di Torino, Italy

- Participation to MIMOSA – MISure sperimentali e MOdelli di traffico dati multiServizio a pAcchetto (Experimental Measurements and Traffic Models of Multiservice Packet Networks), MIUR PRIN 2005-2007, Prof. A. Pattavina, Politecnico di Milano, Italy
- Participation to PATTERN – Protocols And cross-layer Techniques To achieve Energy-efficiency in Reconfigurable adhoc and sensor Networks, MIUR PRIN 2003-2005, Prof. G. Mazzini, Università di Ferrara, Italy
- Participation to PRIMO – Piattaforme Riconfigurabili per Interoperabilità in MOBilità (Reconfigurable Platforms for Mobile Interoperability), MIUR FIRB 2002-2005, Prof. S. Benedetto, Politecnico di Torino, Italy.

Publications

Performance indicators

As of early 2014, my H-index is 17, my i10-index is 24, and my publications have a total of around 1300 citations, according to Google Scholar.

Co-author list

- Silvia Ancona, Politecnico di Bari / INSA Lyon
- José M. Barcelo-Ordinas, Universitat Politècnica de Catalunya
- Christian Bonnet, Eurecom
- Carlo Borgiattino, Politecnico di Torino
- María Calderón Pastor, Universidad Carlos III de Madrid
- Aline Carneiro Viana, Inria
- Claudio Casetti, Politecnico di Torino
- Carla-Fabiana Chiasserini, Politecnico di Torino
- Michele Garetto, Università di Torino
- Tasos Giannoulis, Rice University
- Marco Gramaglia, Istituto Superiore Mario Boella
- Jérôme Härri, Eurecom
- Katia Jaffres-Runser, Université de Toulouse
- Edward Knightly, Rice University
- Chi-Anh La, Eurecom
- Fred Le Moul, INSA Lyon
- Francesco Malandrino, Politecnico di Torino
- Pietro Michiardi, Eurecom
- Diala Naboulsi, INSA Lyon / Inria
- Alessandro Nordio, CNR-IEIIT
- Panos Papadimitratos, KTH

- Massimo Reineri, Istituto Superiore Mario Boella
- Pedro Olmo Stancioli Vaz de Melo, Universidade Federal de Minas Gerais
- Razvan Stanica, ENSEEIHT and INSA Lyon
- Oscar Trullols, Universitat Politècnica de Catalunya
- Sandesh Uppoor, INSA Lyon / Inria

Book chapters

- S. Uppoor, M. Fiore, J. Härri, *Synthetic mobility traces for vehicular networking*, in H. Labiod and A.-L. Beylot (editors), *Vehicular Networks: Models and Algorithms*, Wiley, 2013
- M. Reineri, C. Casetti, C.-F. Chiasserini, M. Fiore, O. Trullols-Cruces, J.M. Barcelo-Ordinas, *RSU Deployment for Content Dissemination and Downloading in Intelligent Transportation Systems*, in R. Daher and A. Vinel (editors), *Roadside Networks for Vehicular Communications: Architectures, Applications, and Test Fields*, Information Science Publishing, 2012
- M. Fiore, C. Casetti, C.-F. Chiasserini, *Information Sharing in VANETs*, in M. Wafta (editor), *Advances in Vehicular Ad-Hoc Networks: Developments and Challenges*, Information Science Publishing, 2010
- M. Fiore, *Vehicular Mobility Models*, in S. Olariu and M. Weigle (editors), *Vehicular Networks from Theory to Practice*, Chapman & Hall/CRC, 2009
- V. Baiamonte, M. Fiore, C. Casetti, C.-F. Chiasserini, *Cross-layer Solutions for Traffic Forwarding in Mesh Networks*, in E. Hossain and K.K. Leung (editors), *Wireless Mesh Networks: Architectures and Protocols*, Springer-Verlag, 2007

Peer-reviewed international journals

- F. Malandrino, C. Borgiattino, C. Casetti, C.-F. Chiasserini, M. Fiore, R. Sadao, *Verification and Inference of Positions in Vehicular Networks through Anonymous Beaconing*, *IEEE Transactions on Mobile Computing*, to appear
- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, *Content Download in Vehicular Networks in Presence of Noisy Mobility Prediction*, *IEEE Transactions on Mobile Computing*, to appear
- S. Uppoor, O. Trullols-Cruces, M. Fiore, J.M. Barcelo-Ordinas, *Generation and Analysis of a Large-scale Urban Vehicular Mobility Dataset*, *IEEE Transactions on Mobile Computing*, to appear
- M. Fiore, C. Casetti, C.-F. Chiasserini, D. Borsetti, *Persistent Localized Broadcasting in VANETs*, *IEEE Journal on Selected Areas in Communications*, to appear
- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, *Optimal Content Downloading in Vehicular Networks*, *IEEE Transactions on Mobile Computing*, Vol. 12, No.7, July 2013
- M. Fiore, C. Casetti, C.-F. Chiasserini, P. Papadimitratos, *Discovery and Verification of Neighbor Positions in Mobile Ad Hoc Networks*, *IEEE Transactions on Mobile Computing*, Vol. 12, No.2, January 2013

- C.-A. La, P. Michiardi, C. Casetti, C.-F. Chiasserini, M. Fiore, *Content Replication in Mobile Networks*, IEEE Journal on Selected Areas in Communications, Vol.30, No.9, October 2012
- O. Trullols-Cruces, M. Fiore, J.M. Barcelo-Ordinas, *Cooperative download in vehicular environments*, IEEE Transactions on Mobile Computing, Vol.11, No.4, April 2012
- S. Uppoor, M. Fiore, *Vehicular mobility in large-scale urban environments*, ACM Mobile Computing and Communications Review, Vol.15, No.4, October 2011
- D. Borsetti, M. Fiore, C. Casetti, C.-F. Chiasserini, *An Application-level Framework for Information Dissemination and Collection in Vehicular Networks*, Elsevier Performance Evaluation, Vol.68, No.9, September 2011
- M. Fiore, C. Casetti, C.-F. Chiasserini, *Caching Strategies Based on Information Density Estimation in Wireless Ad Hoc Networks*, IEEE Transactions on Vehicular Technology, Vol.60, No.15, June 2011
- J. Härrri, M. Fiore, F. Filali, C. Bonnet, *Vehicular Mobility Simulation with VanetMobiSim*, Transactions of The Society for Modeling and Simulation, Vol.87, No.4, April 2011
- O. Trullols-Cruces, J.M. Barcelo-Ordinas, M. Fiore, *Exact Decoding Probability Under Random Linear Network Coding*, IEEE Communication Letters, Vol.15, No.1, January 2011
- O. Trullols-Cruces, M. Fiore, C. Casetti, C.-F. Chiasserini, J.M. Barcelo-Ordinas, *Planning Roadside Infrastructure for Information Dissemination in Intelligent Transportation Systems*, Elsevier Computer Communications, Vol.33, No.4, March 2010
- M. Fiore, C. Casetti, C.-F. Chiasserini, *Information Density Estimation for Content Retrieval in MANETs*, IEEE Transactions on Mobile Computing, Vol.8, No.3, March 2009
- M. Fiore, C. Casetti, G. Galante, *Concurrent Multipath Communication for Real Time Traffic*, Elsevier Computer Communications, Vol.30, No.17, November 2007
- M. Fiore, C. Casetti, C.-F. Chiasserini, M. Garetto, *Analysis and Simulation of a Content Delivery Application for Vehicular Wireless Networks*, Elsevier Performance Evaluation, Vol.64, No.5, June 2007.

Peer-reviewed international conference proceedings

- M. Gramaglia, O. Trullols-Cruces, D. Naboulsi, M. Fiore, M. Calderon, *Vehicular Networks on two Madrid Highways* IEEE SECON 2014, Singapore, June 2014
- S. Ancona, R. Stanica, M. Fiore, *Performance Boundaries of Massive Floating Car Data Offloading*, Invited paper, WONS 2014, Obergurgl, Austria, April 2014

- D. Naboulsi, R. Stanica, M. Fiore, *Classifying Call Profiles in Large-scale Mobile Traffic Datasets*, IEEE INFOCOM 2014, Toronto, Canada, April 2014
- P. Vaz de Melo, A. Viana, M. Fiore, K. Jaffrès-Runser, F. Le Mouel, A. Loureiro *RECAST: Telling Apart Social and Random Relationships in Dynamic Networks*, ACM/IEEE MSWiM 2013, Barcelona, Spain, September 2013
- D. Naboulsi, M. Fiore *On the Instantaneous Topology of a Large-scale Urban Vehicular Network: the Cologne case*, ACM MobiHoc 2013, Bangalore, India, July 2013
- O. Trullols-Cruces, M. Fiore, J.M. Barcelo-Ordinas *Understanding, Modeling and Taming Mobile Malware Epidemics in a Large-scale Vehicular Network*, IEEE WoWMoM 2013, Madrid, Spain, June 2013
- R. Stanica, M. Fiore, F. Malandrino *Offloading Floating Car Data*, IEEE WoWMoM 2013, Madrid, Spain, June 2013
- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, R.S. Yokoyama, C. Borgiattino *A-VIP: Anonymous Verification and Inference of Positions in Vehicular Networks*, IEEE INFOCOM 2013 Miniconference, Turin, Italy, April 2013
- M. Fiore, A. Nordio, C.-F. Chiasserini *Investigating the Accuracy of Mobile Urban Sensing*, WONS 2013, Banff, AB, Canada, March 2013
- S. Uppoor, M. Fiore, *Insights on metropolitan-scale vehicular mobility from a networking perspective*, Invited paper, ACM HotPlanet 2012, Low Wood Bay, UK, June 2012
- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, *Offloading Cellular Networks through ITS Content Download*, IEEE SECON 2012, Seoul, Korea, June 2012
- S. Uppoor, M. Fiore, *Large-scale Urban Vehicular Mobility for Networking Research*, IEEE VNC 2011, Amsterdam, The Netherlands, November 2011
- M. Fiore, C. Casetti, C.-F. Chiasserini, P. Papadimitratos, *Secure Neighbor Position Discovery in Vehicular Networks*, IEEE/IFIP MedHocNet 2011, Favignana, Italy, June 2011
- L. Garelli, C. Casetti, C.-F. Chiasserini, M. Fiore, *MobSampling: V2V Communications for Traffic Density Estimation*, IEEE VTC Spring 2011, Budapest, Hungary, May 2011
- F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, *Content Downloading in Vehicular Networks: What Really Matters*, IEEE INFOCOM 2011 Miniconference, Shanghai, China, April 2011
- C.-A. La, P. Michiardi, C. Casetti, C.-F. Chiasserini, M. Fiore, *A Lightweight Distributed Solution to Content Replication in Mobile Networks*, IEEE WCNC 2010, Sydney, Australia, April 2010
- D. Borsetti, M. Fiore, C. Casetti, C.-F. Chiasserini, J.M. Barcelo-Ordinas, *Virtual Data Mules for Data Collection in Road-side Sensor Networks*, ACM MobiOpp 2010, Pisa, Italy, February 2010

- D. Borsetti, M. Fiore, C. Casetti, C.-F. Chiasserini, *Cooperative Support for Localized Services in VANETs*, ACM MSWiM 2009, Tenerife, Spain, October 2009
- O. Trullols-Cruces, M. Fiore, C. Casetti, C.-F. Chiasserini, J.M. Barcelo-Ordinas, *A Max Coverage Formulation for Information Dissemination in Vehicular Networks*, IEEE WiMob 2009, Marrakech, Morocco, October 2009
- M. Fiore, J.M. Barcelo-Ordinas, *Cooperative download in urban vehicular networks*, IEEE MASS 2009, Macau, China, October 2009
- C. Casetti, C.-F. Chiasserini, M. Fiore, C.-A. La, P. Michiardi, *P2P Cache-and-Forward Mechanisms for Mobile Ad Hoc Networks*, IEEE ISCC 2009, Sousse, Tunisia, July 2009
- M. Fiore, F. Mininni, C. Casetti, C.-F. Chiasserini, *To Cache or Not To Cache?*, IEEE INFOCOM 2009, Rio de Janeiro, Brazil, April 19-25, 2009
- A. Giannoulis, M. Fiore, E. Knightly, *Supporting Vehicular Mobility in Urban Multi-hop Wireless Networks*, ACM/USENIX MobiSys 2008, Breckenridge, CO, USA, June 17-20, 2008
- M. Fiore, J. Härri, *The Networking Shape of Vehicular Mobility*, ACM MobiHoc 2008, Hong Kong, China, May 2008
- M. Fiore, J. Härri, F. Filali, C. Bonnet, *Understanding Vehicular Mobility in Network Simulation*, IEEE MoVeNet 2007, Pisa, Italy, October 2007
- C. Casetti, C.-F. Chiasserini, M. Fiore, *Relay quality awareness in mesh networks routing*, Tyrrhenian International Workshop on Digital Communication, Ischia, Napoli, Italy, September 2007
- M. Fiore, C. Casetti, C.-F. Chiasserini, *Efficient Retrieval of User Contents in MANETs*, IEEE INFOCOM 2007, Anchorage, AK, USA, May 2007
- M. Fiore, J. Härri, Fethi Filali, Christian Bonnet, *Vehicular Mobility Simulation for VANETs*, SCS/IEEE Annual Simulation Symposium, Norfolk, VA, USA, March 2007
- C. Rossi, C. Casetti, M. Fiore, D. Schonfeld, *A Partially Reliable Transport Protocol For Multiple-Description Real-Time Multimedia Traffic*, IEEE ICIP 2006, Atlanta, GA, USA, October 2006
- M. Fiore, C. Casetti, *An Adaptive Transport Protocol for Balanced Multihoming of Real-Time Traffic*, IEEE GLOBECOM 2005, St.Louis, MO, USA, November 2005
- M. Fiore, C. Casetti, C.-F. Chiasserini, *On-demand Content Delivery in Vehicular Wireless Networks*, IEEE/ACM MSWIM 2005, Montreal, Canada, October 2005
- C. Casetti, C.-F. Chiasserini, M. Fiore, M. Garetto, *Notes on the Inefficiency of 802.11e HCCA*, IEEE VTC 2005-Fall, Dallas, TX, USA, September 2005.

Posters, short papers, announcements

- D. Naboulsi, M. Fiore, C.-F. Chiasserini *Assessment of Practical Energy Savings in Cellular Networks*, IEEE INFOCOM 2014, Toronto, Canada, April 2014
- D. Naboulsi, M. Fiore, R. Stanica *On the Characterization of Mobile Calling Behaviors*, ACM MobiHoc 2013, Bangalore, India, July 2013
- D. Naboulsi, M. Fiore, R. Stanica *Human Mobility Flows in the City of Abidjan*, NetMob 2013, Boston, MA, USA, May 2013
- D. Naboulsi, M. Fiore, *The Connectivity of Colognes Large-scale Vehicular Network*, IEEE INFOCOM 2013, Torino, Italy, April 2013
- S. Uppoor, M. Fiore, *Vehicular mobility in large-scale urban environments*, ACM MobiCom 2011, Las Vegas, NV, USA, September 2011
- P. Michiardi, C.-F. Chiasserini, C. Casetti, C.-A. La, M. Fiore, *On a Selfish Caching Game*, ACM PODC 2009, Calgary, Canada, August 2009
- J. Härri, M. Fiore, F. Filali, C. Bonnet, *VanetMobiSim: Generating Realistic Mobility Patterns for VANETs*, ACM VANET 2006, Los Angeles, CA, USA, September 2006
- M. Fiore, C. Casetti, *Multihoming for real-time traffic: the Westwood SCTP-PR protocol*, IEEE INFOCOM 2005 - Student Workshop Poster Session, Miami, FL, USA, March 2005.

Theses

- M. Fiore, *Vehicular Networking: simulation, applications, experimentation*, PhD thesis, Politecnico di Torino, April 2008
- M. Fiore, *Quality of service support in IEEE 802.11e wireless LANs*, MS thesis, University of Illinois at Chicago and Politecnico di Torino, July 2004.

Professional certifications

Feb 2009	French qualification to university faculty positions
Nov2004	Italian professional qualification in Computer Science engineering
Jun 2002	ETS <i>Test of English as a Foreign Language</i> (TOEFL), score 273/300.

Affiliations

Member of IEEE.

Awards

- French national Research Excellence Award (*Prime pour l'Excellence Scientifique*), October 2012
- PhD scholarship from *Istituto Superiore Mario Boella*, funding research activities on wireless networks, December 2004.

