Understanding and Exploiting Mobility in Wireless Networks

Sandesh Uppoor

Number 2013-ISAL-0103 Year 2013

THÈSE

UNDERSTANDING AND EXPLOITING MOBILITY IN WIRELESS NETWORKS

pour obtenir le grade de docteur délivreè par

L'INSTITUT NATIONAL DES SCIENCES APPLIQUÈES DE LYON

présentée et soutenue publiquement par

Mr Sandesh Uppoor

le 29 / 11 / 2013

Directeur de thèse **Prof. Fabrice Valois**
Co-encadrement de la thèse **Prof. Marco Fiore**

Jury

Prof. André-Luc BEYLOT	IRIT, ENSEEIHT, France
Dr. Artur ZIVIANI	LNCC, Brazil
Prof. Claudio CASETTI	Politecnico di Torino, Italy
Dr. Laurent ROULLET	Alcatel-Lucent Bell Labs, France
Prof. Yacine GHAMRI-DOUDANE	ENSIIE, France
Dr. Marco FIORE	INSA-Lyon, France
Prof. Fabrice VALOIS	INSA-Lyon, France

Number 2013-ISAL-0103 Year 2013

THESIS

UNDERSTANDING AND EXPLOITING MOBILITY IN WIRELESS NETWORKS

Dissertation submitted to

L'INSTITUT NATIONAL DES SCIENCES APPLIQUÈES DE LYON

in accordance with the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

on 29 / 11 / 2013

by

Mr Sandesh Uppoor

Supervisor **Prof. Fabrice Valois**

Co-Supervisor **Prof. Marco Fiore**

Jury

Prof. André-Luc BEYLOT	IRIT, ENSEEIHT, France
Dr. Artur ZIVIANI	LNCC, Brazil
Prof. Claudio CASETTI	Politecnico di Torino, Italy
Dr. Laurent ROULLET	Alcatel-Lucent Bell Labs, France
Prof. Yacine GHAMRI-DOUDANE	ENSIIE, France
Dr. Marco FIORE	INSA-Lyon, France
Prof. Fabrice VALOIS	INSA-Lyon, France

This work is carried out in CITI Laboratory and was supported by the joint laboratory between INRIA and Alcatel-Lucent Bell Labs on Self Organizing Networks.

Acknowledgements

First and the foremost, I would like to extend my sincere thanks to my advisors Dr. Marco Fiore and Prof. Fabrice Valois for their support and trust. During all these three years, they have consistently helped me to improve my weak aspects in research. I wholeheartedly appreciate their time and patience which they have invested that made this thesis work a success. I would also like to thank Dr. Laurent Roullet for his valuable assistance.

I take this opportunity to thank the jury members for making the effort to read and review this thesis, for taking a long trip to attend the defence, and for giving me useful feedback.

Days in CITI laboratory and in Lyon have marked the most memorable part in my life, I like to thank my dear friends and colleagues for good, friendly, and professional, atmosphere all these years. I specially would like to thank IT and administrative staff at CITI for their instant support whenever I asked for one.

Finally, I would like to extend my special gratitude to my family and friends back home in India. Their distant support has always helped me to bounce back whenever I felt low.

Résumé

Le degré de pénétration du marché des appareils intelligents tels que les smartphones et les tablettes avec les technologies de communication embarquées comme le WiFi, 3G et LTE a explosé en moins d'une décennie. En complément de cette tendance technologique, les applications des réseaux sociaux ont virtuellement connecté une grande partie de la population, en génèrant une demande de trafic de données croissant vers et depuis l'infrastructure de communication. Les communications pervasive ont aussi acquis une importance dans l'industrie automobile. L'émergence d' une gamme impressionnante d' appareils intelligents dans les véhicules permettant services tels que assistance au conducteur, infotainment, suivi à distance du vehicule, et connectivité àux réseaux sociaux même en déplacement. La demande exponentielle de connectivité a encore défié les fournisseurs de services de télécommunications pour répondre aux attentes des utilisateurs du réseau à grande vitesse.

L'objectif de cette thèse est de modéliser et comprendre la mobilité dynamique des utilisateurs à grande vitesse et leurs effets sur les architectures de réseau sans fil.

Compte tenu de l' importance du développement de notre étude sur une représentation réaliste de la mobilité des véhicules, nous étudions tout d'abord les approches les plus populaires pour la génération de trafic routier synthétique et discutons les caractéristiques des ensembles de données accessibles au public qui decrivent des mobilités véhiculaires. En utilisant l'information des déplacements de la population dans une région métropolitaine, les données du réseau routier détaillées et des modèles réalistes de conduite microscopiques, nous proposons un jeux de données de mobilité véhiculaire original qui redéfinit l'état de l'art et qui replie la circulation routière de facon realiste dans le temps et dans l'espace. Nous étudions ensuite l'impact des dynamiques de mobilité du point de vue de la couverture cellulaire en présence d'un déploiement réel des stations de base. En outre, en examinant les effets de la mobilité des véhicules sur les réseaux autonomes, nous voyons des possibilités pour les futurs paradigmes de réseaux hétérogènes.

Motivés par l'évolution dynamique dans le temps de la mobilité des véhicules observée dans notre jeux de données, nous proposons également une approche en ligne pour prédire les flux de trafic macroscopiques. Nous analysons les paramètres affectant la prédiction de la mobilité en milieu urbain. Nous dévoilons *quand* et *où* la gestion des ressources réseau est plus crucial pour accueillir le trafic généré par les utilisateurs à bord. Ces études dévoilent des multiples opportunités de gestion intelligente des transports, soit pour construire de nouvelles routes, soit pour l'installation de bornes de recharge électriques, ou pour la conception de systèmes de feux de circulation intelligents, contribuant ainsi à la planification urbaine.

Abstract

The market penetration of smart devices like smartphones and tablets with embedded communication technologies like WiFi, 3G and LTE has exploded in less than a decade. Complementing this technological trend, social networking applications have virtually connected a large portion of the population generating an ever-growing data traffic demand on the communication infrastructure. Pervasive communications have gained significance in the automobile industry as well, with the emergence of an impressive range of in-vehicle smart devices enabling driver assistance, infotainment, over-the-air vehicle monitoring, and even social connectivity on the move. This surge in the demand for connectivity has further challenged telecommunication service providers to meet the expectations of high-speed network users.

The goal of this thesis is to model and understand the mobility dynamics of high-speed users and their effect on wireless network architectures.

Given the importance of developing our study on a realistic representation of vehicular mobility, we first survey the most popular approaches for the generation of synthetic road traffic and discuss the features of publicly available vehicular mobility datasets. Using original travel demand information of the population of a metropolitan area, detailed road network data and realistic microscopic driving models, we propose a novel state-of-art vehicular mobility dataset that closely mimics the real-world road traffic dynamics in both time and space. We then study the impact of such mobility dynamics from the perspective of wireless cellular network architecture in presence of a real-world base station deployment. In addition, by discussing the effects of vehicular mobility on autonomous network architecture, we hint at the opportunities for future heterogenous network paradigms.

Motivated by the time-evolving mobility dynamics observed in our original dataset, we also propose an online approach to predict near-future macroscopic traffic flows. We analyze the parameters affecting the mobility prediction in an urban environment and unveil *when* and *where* network resource management is more crucial to accommodate the traffic generated by users onboard. Such studies unveil multiple opportunities in transportation management either for building new roads, installing electric charging points, or for designing intelligent traffic light systems, thereby contributing to urban planning.

Contents

Acknowledgements iii

Résumé v

Abstract vii

List of Figures xiv

List of Tables xv

1 Introduction 1
- 1.1 Context . . . 2
- 1.2 Cellular access technologies . . . 3
- 1.3 Autonomous network technologies . . . 7
- 1.4 Contribution and outline . . . 8

2 Synthetic vehicular mobility 11
- 2.1 Introduction . . . 12
- 2.2 Real-world vehicular mobility datasets . . . 12
- 2.3 Generation process . . . 14
  - 2.3.1 Road topology database . . . 14
  - 2.3.2 Microscopic traffic flow description . . . 16
  - 2.3.3 Macroscopic road traffic description . . . 18
- 2.4 Mobility simulators . . . 20
  - 2.4.1 Microscopic traffic simulators . . . 20
  - 2.4.2 Mesoscopic traffic simulators . . . 21
  - 2.4.3 Macroscopic traffic simulators . . . 22
  - 2.4.4 Interactions between simulators . . . 22
- 2.5 Mobility datasets . . . 24
  - 2.5.1 Perception . . . 24
  - 2.5.2 Small-scale measurements . . . 28
  - 2.5.3 Road traffic imagery . . . 29
  - 2.5.4 Roadside detectors . . . 30
  - 2.5.5 Socio-demographic surveys . . . 31
  - 2.5.6 Discussion . . . 32
- 2.6 Summary . . . 33

3 A large-scale realistic vehicular mobility scenario: Koln, Germany 35
- 3.1 Introduction . . . 36
- 3.2 Dataset generation process . . . 36
  - 3.2.1 Road topology . . . 36

3.2.2 Microscopic vehicular mobility . . . 37
3.2.3 Traffic demand . . . 37
3.2.4 Traffic assignment . . . 37
3.2.5 Simulation . . . 38
3.3 Repairing the dataset . . . 39
3.3.1 Over-comprehensive and bursty traffic demand . . . 40
3.3.2 Inconsistent road information . . . 42
3.3.3 Flawed road topology conversion . . . 44
3.3.4 Simplistic default traffic assignment . . . 46
3.4 Koln vehicular mobility dataset . . . 47
3.5 Integrating the Koln dataset into OpenAirInterface . . . 53
3.6 Summary . . . 55

4 Impact of user mobility in wireless networks 57
4.1 Introduction . . . 58
4.2 Access network . . . 58
4.2.1 RAN infrastructure deployment . . . 58
4.2.2 Macroscopic analysis . . . 59
4.2.3 Microscopic analysis . . . 68
4.3 Autonomous network . . . 77
4.3.1 Modified Koln dataset . . . 77
4.3.2 Vehicular contact . . . 78
4.3.3 Network graph metrics . . . 79
4.3.4 Epidemic dissemination . . . 84
4.4 Summary . . . 87

5 Online macroscopic vehicular mobility prediction 89
5.1 Introduction . . . 90
5.2 Related work . . . 91
5.2.1 Basic prediction approach . . . 91
5.2.2 Vehicular mobility prediction on road networks . . . 92
5.3 Macroscopic vehicular mobility prediction models . . . 94
5.3.1 Markovian model . . . 94
5.3.2 Sequence Matching model . . . 98
5.3.3 Online prediction with fallback . . . 99
5.3.4 Estimating the prediction accuracy . . . 100
5.4 Scenarios . . . 102
5.4.1 Vehicular mobility adaptation for prediction model. . . . 103
5.4.2 Impact of model parameters for crossroads . . . 104
5.4.3 Impact of model parameters for the cellular network scenario . . . 107
5.5 Performance evaluation . . . 110
5.5.1 Road network scenario . . . 110
5.5.2 Cellular network . . . 120
5.6 Summary . . . 126

6 Conclusions and perspectives 129
6.1 Conclusion . . . 129
6.2 Perspectives . . . 132

Contents

List of publications 133

Bibliography 135

List of Figures

1.1 Market penetration and forecast of smart devices (Source: International Data Corporation) . . . 3
1.2 Mobile Broadband users (Source: ITU World Telecommunication /ICT Indicators database) . . . 4
1.3 User mobility versus data rates. . . . 5
1.4 Vehicular user scenario and applications (Source: [CSZ11]). . . . 6

2.1 Concept map for synthetic vehicular datasets, where the three main blocks (dotted blocks) are enriched with added blocks and features providing an increased accuracy. . . . 15
2.2 Road topology database examples. . . . 16
2.3 Traffic demand and traffic assignment. . . . 19
2.4 The microscopic simulator SUMO calibrated by external simulators. . . . 23
2.5 Interactions between network and traffic simulators . . . 24
2.6 Turin scenario: downtown, sparse traffic. This figure is best viewed in colors. 26
2.7 Turin scenario: city outskirts, dense traffic. This figure is best viewed in colors. 27
2.8 Zurich scenario: downtown. This figure is best viewed in colors. . . . 28
2.9 Karlsruhe scenario: downtown. This figure is best viewed in colors. . . . 29
2.10 Zurich scenario: region (city and surroundings). This figure is best viewed in colors. . . . 31

3.1 Simulation workflow. . . . 38
3.2 Original TAPASCologne dataset: Traffic features over time. . . . 39
3.3 Original TAPASCologne dataset. Snapshot of the traffic status at 7:00 am, in a 400 km^2 region centered on the city of Koln. This figure is best viewed in colors. . . . 40
3.4 Volume of traffic injected in the road network between 6:00 am and 12:00 pm, according to the original and fixed TAPASCologne O/D matrix . . . 41
3.5 Example of wrong restriction enforcement in OpenStreetMap data . . . 42
3.6 Example of continuous restriction enforcement in OpenStreetMap . . . 43
3.7 Example of unrecognized information ignored during map conversion . . . 44
3.8 Example of topological information unfitness during map conversion . . . 45
3.9 Traffic evolution over multiple iterations of the assignment algorithm . . . 46
3.10 Final TAPASCologne dataset. Traffic features over time . . . 47
3.11 Spatiotemporal distribution of vehicular traffic density (5 am to 10 am). . . . 48
3.11 Spatiotemporal distribution of vehicular traffic density (11 am to 4 pm). . . . 49
3.11 Spatiotemporal distribution of vehicular traffic density (5 pm to 10 pm). . . . 50
3.12 Final TAPASCologne dataset. Trip feature distributions . . . 51

3.13 Final TAPASCologne dataset. Traffic at 5:00 pm, in the real world (left) and in simulation (right). The plots refer to the whole urban area (top) and in the city center (bottom). This figure is best viewed in colors. 53
3.14 Static mobility feed from SUMO to OpenAirLTE. 54
3.15 Real time mobility feed from SUMO to OpenAirLTE through TraCI. 55

4.1 Deployment of the RAN infrastructure in Koln, for one mobile operator. Base station locations (left), and cell layout resulting from the Voronoi tessellation (right). 59
4.2 (a)–(i): spatiotemporal evolution of the expected data traffic load generated by pervasive vehicular access in the Koln region, during a typical day (5 am to 1 pm). 60
4.2 (j)–(r): spatiotemporal evolution of the expected data traffic load generated by pervasive vehicular access in the Koln region, during a typical day (2 pm to 10 pm). 61
4.3 (a): typical daily profile of normalized RAN traffic load – concession of the Autonomous Networks Research Group at USC, `http://anrg.usc.edu`. (b): geographical distribution of the population in the Koln region. 62
4.4 Macroscopic vehicular flows in the Koln dataset. Plots refer to four different hours of the day. This figure need to be viewed in colors. 64
4.5 Spatiotemporal evolution of macroscopic flows of pervasive vehicular access in the Koln region over one day. This figure is best viewed in colors (5 am to 1 pm). 66
4.5 Spatiotemporal evolution of macroscopic flows of pervasive vehicular access in the Koln region over one day. This figure is best viewed in colors (2 pm to 10 pm). 67
4.6 Distributions of vehicular user inter-arrival time at RAN cells. 69
4.7 Spatiotemporal evolution of inter-arrival times (in seconds) at cells deployed in the Koln region during over one day (5 am to 1 pm). 70
4.7 Spatiotemporal evolution of inter-arrival times (in seconds) at cells deployed in the Koln region during over one day (2 pm to 10 pm). 71
4.8 Theoretical distribution fittings on inter-arrival times in the Koln RAN deployment for different hours of a typical day. Each fitting sample is shown in linear-linear and logarithmic-logarithmic scale plots. 72
4.9 Distributions of vehicular user residence time at RAN cells. 74
4.10 Theoretical distribution fittings of cell residence times in the Koln cellular deployment for different hours of a typical day. Each fitting sample is shown in linear-linear and logarithmic-logarithmic plots. 75
4.11 Koln pruned scenario: city and surroundings. 78
4.12 Koln scenario: downtown. 79
4.13 Distributions of the inter-vehicle contact duration, over the whole day. 80
4.14 Average number of clusters, with singletons, and mean cluster size. 81
4.15 *Region* scenarios. Clusters in time (left) and cluster size CDF (right). 81
4.16 *Downtown* scenarios. Cluster size (left) and degree (right) CDF. 82
4.17 Epidemic dissemination ratio over time with a penetration rate of 1.0 (left) and latency to achieve the dissemination ratio quantiles 0.05, 0.25, 0.5, 0.75 and 0.95 under different penetration rates (right). 83

4.18 Survivability of the epidemic dissemination. Penetration rate over the whole day (left) and during the morning rush hours (right) 85
4.19 Initial spreading of the epidemic dissemination. Evolution of the dissemination ratio over time when the penetration rate is 1.0 (left) and latency to achieve the dissemination ratio quantiles 0.05, 0.25, 0.5, 0.75 and 0.95 in presence of different penetration rates (right). 86

5.1 Different phases in basic prediction process. 91
5.2 Sample network representation. 94
5.3 State transition diagram of order-1 and a subset of the order-2 Markovian models. 95
5.4 q-step prediction representation. 97
5.5 Example of sequence matching. 98
5.6 Sliding window of training and observation. 100
5.7 Simple order-1 showing state transition with different traffic behaviors. . . . 101
5.8 Road network and cellular network representation. 102
5.9 Road network: Varying training interval T_B 104
5.10 Road network: Varying observation intervals. 105
5.11 Road network: Overall error of 60 minutes training with different Markovian order. 106
5.12 Road network: Error vs volume-per-state/h over 24 hours calculated for every state. 107
5.13 Cellular network: Varying training intervals. 108
5.14 Cellular network: Varying prediction intervals. 109
5.15 Cellular network: Overall error with 20 minutes training and observation intervals. 109
5.16 Road network: q-step mobility prediction using Markovian model. 111
5.17 Road network: q-step mobility prediction using Sequence Matching. 112
5.18 Road network: Variation of prediction error with time of the day with order-2 Markovian prediction. 113
5.19 Road network: Spatiotemporal distribution of order-2 prediction error and vehicular volume (5 am to 10 am). 114
5.19 Road network: Spatiotemporal distribution of order-2 prediction error and vehicular volume (11 am to 4 pm) 115
5.19 Road network: Spatiotemporal distribution of order-2 prediction error and vehicular volume (5 pm to 10 pm) 116
5.20 Road network: Comparison of 1-step and 5-step prediction results using Markovian model. 117
5.21 Road network: Variation of prediction error with time of the day with order-2 Sequence Matching prediction 118
5.22 Cellular network: q-step mobility prediction using Markovian Model. 120
5.23 Cellular network: q-step mobility prediction using Sequence Matching. 121
5.24 Cellular network : Variation of prediction error with time of the day with order-1 Markovian prediction. 122
5.25 Cellular network: Spatiotemporal distribution of prediction error and vehicular volume (5 am to 10 am). 123
5.25 Cellular network: Spatiotemporal distribution of prediction error and vehicular volume (11 am to 4 pm) 124

5.25 Cellular network: Spatiotemporal distribution of prediction error and vehicular volume (5 pm to 10 pm) 125
5.26 Cellular network : Variation of prediction error with time of the day with order-1 Sequence Matching prediction. 127

List of Tables

2.1 Major features of the real-world vehicular mobility datasets. 14
2.2 Major features of publicly available synthetic vehicular mobility datasets . . . 25

1 Introduction

Contents

1.1 Context 2
1.2 Cellular access technologies 3
1.3 Autonomous network technologies 7
1.4 Contribution and outline 8

1.1 Context

Recent advances in communication technologies have shifted the long-standing paradigm of *we search for information* into *information searches us*. In such a world of seamless connectivity, information is made available on the move to make every step in ones' life better, simpler, fun and more in-touch. The success of smart devices with flexible software platforms have completely changed the way people communicate, and software applications offering Voice over IP (VoIP), instant messaging, photo sharing services have gained unprecedented popularity. According to data analysts, the market penetration rate of smart devices will continue to swell in the near future, as shown in Fig. 1.1. According to actual numbers of 2012, smart devices (i.e., smartphones and tablets) occupies around 70% of the total market share of internet connected devices. It is forecasted that this dominance holds good for next five years. As a result, the demand for high-speed mobile broadband will accelerate in the near future, as illustrated in Fig. 1.2. Referring to the data since 2007, the number of mobile broadband users have increased from 20 to 75 users in developed nations per 100 inhabitants, hence we can speculate a steady growth in the near further as-well looking at the penetration rate of smart devices. This poses a real challenge to telecom service providers in order to maintain the quality of service requested by network users. With the growing urban population, the number of users per square kilometer is also a key concern as the service provider is obliged to efficiently and intelligently manage the available resources to accommodate the generated demand. To address such ever more stringent requirements and to maximize the reusability of available radio spectrum, service providers are committed to the adoption of low-power radio access nodes with a service coverage from 10 meters to 1 kilometer.

Despite such efforts, characterizing the impact of uncontrolled user mobility on the service provided remains an open challenge. That is especially critical in vehicular environments. Incorporation of communication technologies like Ford SYNC, Mercedes-Benz DriveKit, Chevrolet MyLink, Buick IntelliLink, Cadillac CUE in vehicles has unveiled the great service potential of the networked vehicle paradigm. As a result in-vehicle services like turn-by-turn driver assistance, stolen vehicle tracking, real-time traffic and prediction, weather reporting, remote monitoring of vehicle and many more have become one of the major marketing strategies in the automobile industry. Several manufactures have been providing such services via add-on devices and applications like BMW Assist, GM OnStar, Mercedes-Benz mbrace that can be incorporated in vehicles manufactured by other companies that do not offer such services. With such a momentum, these in-vehicle technologies are likely to follow the steps of smart handheld devices in the near future. This further accelerates the challenge to meet the demands of data-hungry high-speed users.

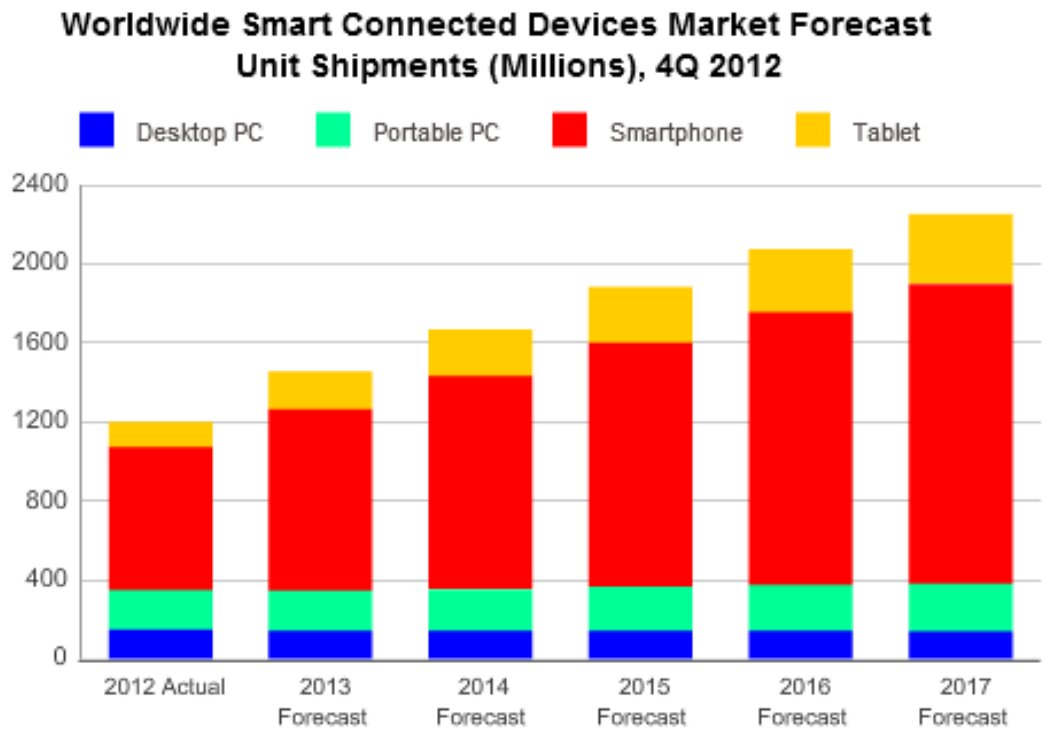

Figure 1.1: Market penetration and forecast of smart devices (Source: International Data Corporation)

1.2 Cellular access technologies

The cellular network is today the most pervasive access network technology with 6.8 billion subscriptions [IUT13]. This number has been growing for the last three decades and this trend is expected to continue in the future. Since 2005 the mobile subscription has seen an increase from 2 to 6.8 billion [IUT13] and is expected to exceed the total world population with such growth rate. This easily indicates the use of multiple subscriptions by an individual. This surge in demand has led to tremendous research interest in this field.

The cellular access network consists of several cells covering a service area. Each cell has a base-station (BS) which serves the users in its vicinity. The users may be either stationary or mobile. Each BS is assigned a specific number of channels, each of which can accommodate one or more calls depending on the transmission and data encoding technology (e.g., FDMA, TDMA, CDMA, OFDMA). One of the major features of a cellular network, in contrast to a traditional public switched telephone network, is user mobility. This implies that when a user moves from one cell to another, the call in progress has to be handed off from one BS to another to ensure continuity of service. If not enough capacity is available in the adjacent

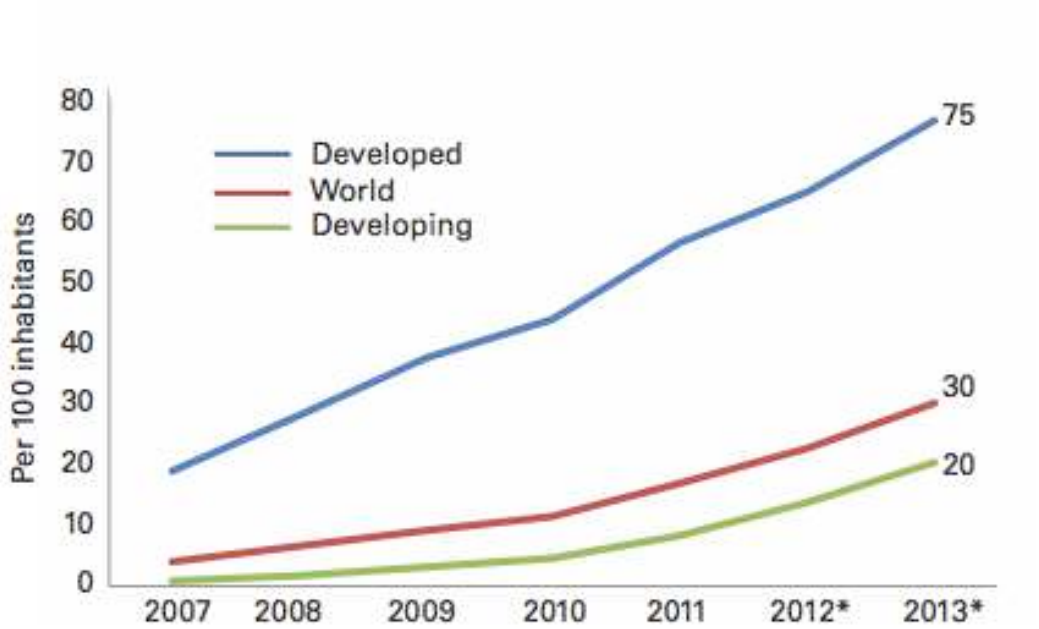

Figure 1.2: Mobile Broadband users (Source: ITU World Telecommunication /ICT Indicators database)

cell, then the call might be interrupted and dropped. Failure to provide uninterrupted service has a direct implication on the popularity of the service provider. Due to the limited capacity assigned to a cell, the handoff decision is influenced by several quality of service (QoS) factors.

The nomenclature of the wireless cellular access generations (G) generally refers to a change in the fundamental nature of the service, non-backward compatible transmission technology, and new frequency bands. In 1981, a transition from analog transmission (1G) to digital transmission (2G) took place. In both cases only voice communication was provided. The introduction of data communication started with *General Packet Radio Service* (2.5G) offering data rates up to 40 Kbps, which allowed a mobile phone to act like any other computer on the internet, sending and receiving data. In order to support the growing demand, communication technologies have constantly been developed to support user' applications that need high data rates. This led to *Enhanced Data rates for GSM Evolution* (2.75G) with a maximum data rate of 384 Kbps. *High Speed Packet Access* and *High-Speed Downlink Packet Access* offer peak data rates of 14 Mbps and 42.2 Mbps respectively, and belong to the third generation (3G) of cellular technologies. *Evolved High Speed Packet Access* (3.5G) supersedes 3G, offering data access at a maximum rate of 168 Mbps. Ultimately, the ubiquitous nature of information access connects today about 2.7 billion mobile phone users, i.e., around 39% [IUT13] of the world population to the Internet.

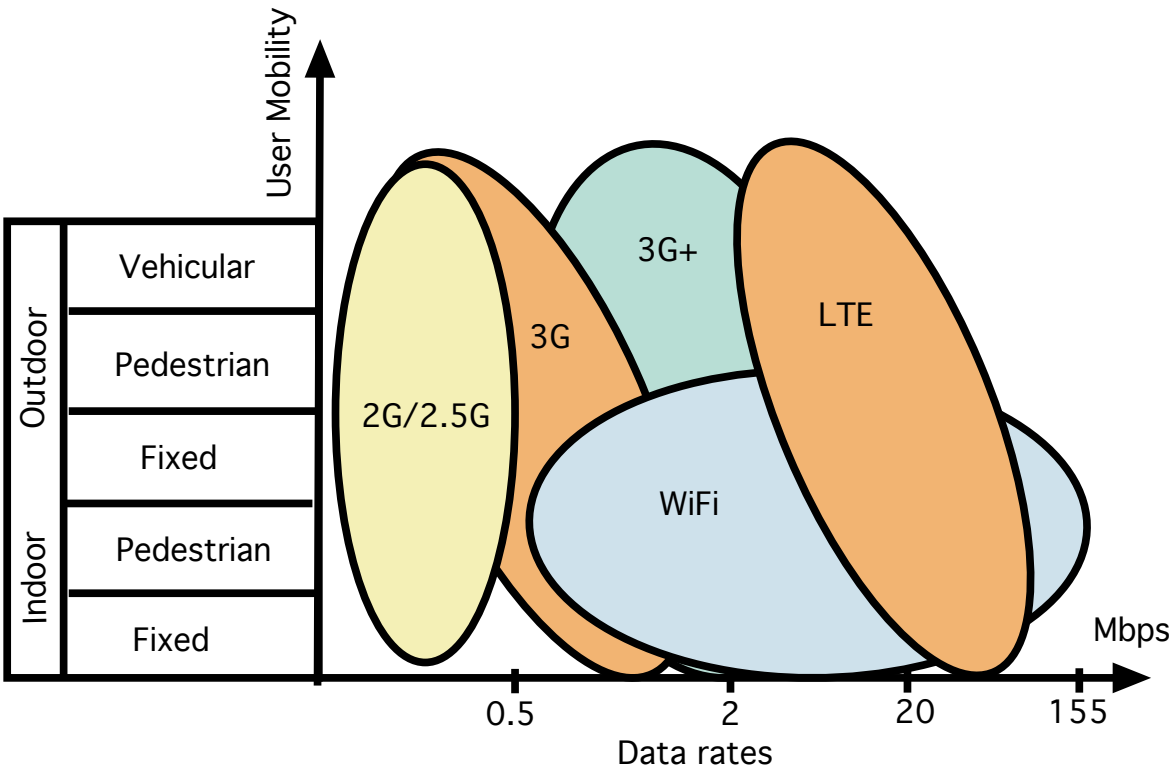

Figure 1.3: User mobility versus data rates.

All these technologies provide data access to users traveling at any mobility profile up to 120 km/h (e.g., a user traveling in a train or car). Fig. 1.3 shows an visual comparison of different technologies, y-axis refers to nature of mobility of the user and x-axis indicate data rates supported by them. We see that all cellular access standards support high-speed users offering increased data rates starting from 2.5G to the recent LTE, whereas in WiFi, better data rates are possible but serves only pedestrian users.

Recently, to meet the growing need of high-speed wireless connections, a new generation of cellular network, i.e., Long Term Evolution (LTE) has been standardized and deployed by the 3rd Generation Partnership Project (3GPP). By employing advanced smart antenna techniques, fast capacity dependent scheduling, adaptive coding and modulation, LTE offers a very high peak data rate of 300 Mbps in ideal conditions. However, the capacity of the LTE network is not evenly distributed, i.e., the cell edge users have much worse throughput than the users near the center [Bra11]. The successor of LTE, the LTE-Advanced, aims at improving the system capacity and the user experience at cell edges. In order to meet the throughput demands and evenly distribute the capacity of a wireless network, a new design paradigm, i.e., the Heterogeneous and Small Cell Networks (HetSNets), was introduced in

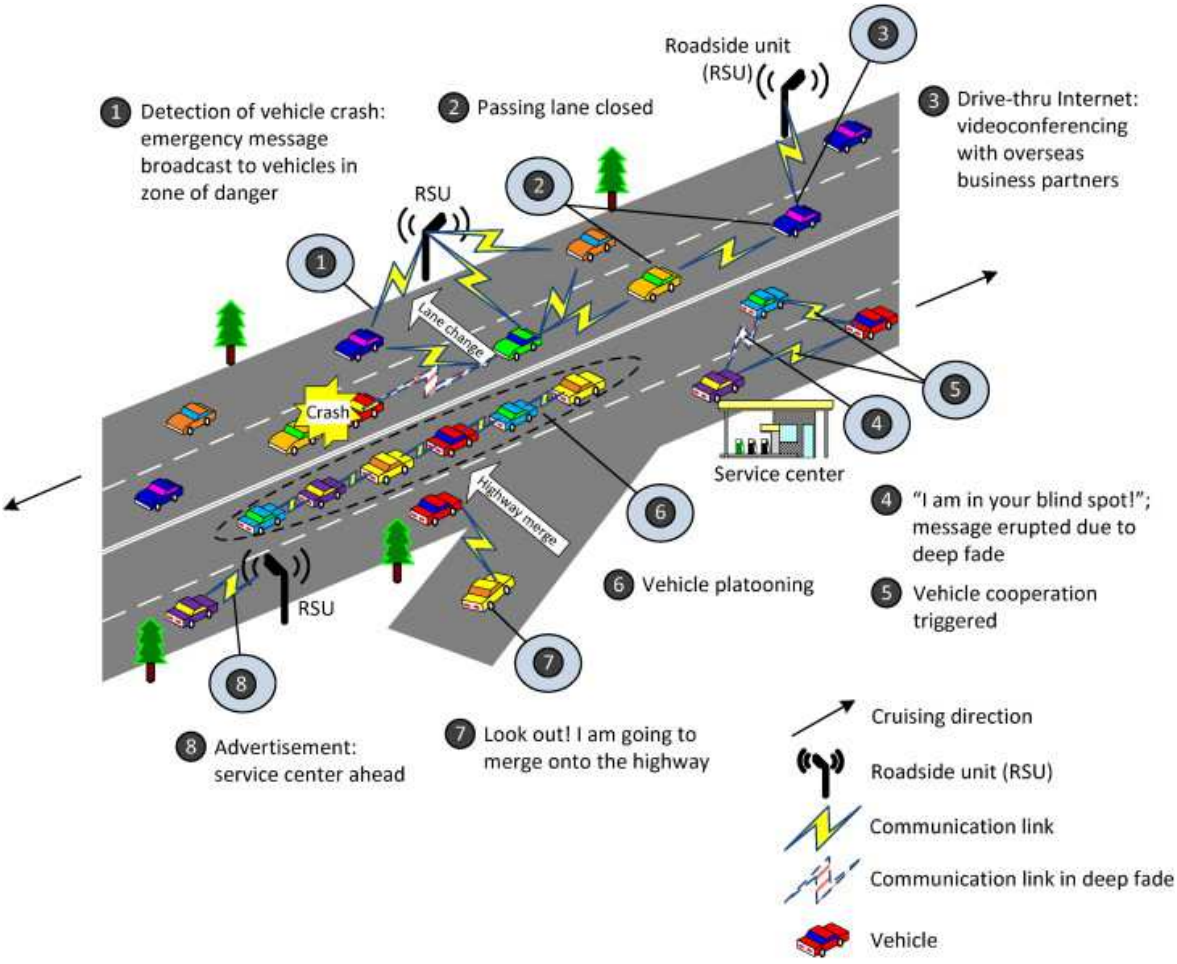

Figure 1.4: Vehicular user scenario and applications (Source: [CSZ11]).

LTE release 10 [DMW+11]. The idea of HetSNets is to deploy several low-power nodes within the coverage of macro BSs to either extend the coverage or boost the local capacity in certain hot-spot areas. This has been the primary strategy for operators to ease the mobile data demand, and is expected to relieve the cellular network from almost 50% of the data traffic by 2017 [Cis13].

One of the effects of such technological advancements is the emergence of networked vehicle applications that will lead us into the age of pervasive vehicular access to radio infrastructure, where vehicles become a main source and destination of large amounts 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, elevated speed and unique movement patterns of vehicular users. In such a scenario, the additional capacity provided by HetSNets may not be used to meet the vehicular user demand because their limited coverage

and the travel velocity of vehicles lead to exceedingly short-lived contacts. As an example, in the case of WiFi offloading, vehicular cell residence time is limited to a mere 13 seconds according to experimental measurements [BHM+06].

Within such a context, it is important to anticipate how pervasive vehicular access will affect the cellular network. Specifically, understanding the unique dynamics of vehicle-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 peak of the data traffic demand originates. Hence, such studies requires accurate vehicular traffic dynamics in a large-scale scenario.

1.3 Autonomous network technologies

Autonomous networks belong to another category of communication networks, where vehicles are not only the source and destination of the information but also spontaneously build the networking system between moving vehicles. Such a concept is implemented by vehicles equipped with wireless interfaces that could have similar or different radio interface technologies, employing short-range to medium-range communication systems. Extensive industrial and academic research has led to the definition of dedicated standards such as IEEE 802.11p [JD08], IEEE 1609 [She13], OSI CALM-M5 [BLJL10] and ETSI ITS [ETS13].

Fig. 1.4 highlights several applications that can be envisioned with vehicular autonomous networks. This concept can facilitate real-life applications by enabling communications among nearby vehicles (vehicle-to-vehicle, V2V) like work zone or accident event notification (1), vehicle lane changing (2), crossroad collision avoidance (7) or grouping vehicles into platoons to increase the capacity of roads (6). These applications best fit the V2V scenario as the events occur instantaneously and are meaningful to a small pool of vehicular participants. Several other applications involving data sharing are of interest to larger audiences: road traffic status, map sharing, location advertisements (8), infotainment and internet access (3) can be best served with communication between vehicles and nearby infrastructure (vehicle-to-infrastructure, V2I) made of road side units (RSU). Such applications enhance the overall driving experiences by avoiding traffic congestion, improving fuel efficiency by smart navigation, travel time management, and most importantly avoiding accident by human error and thereby saving lives.

In comparison with nodes of other autonomous networks like sensor networks and WiFi-enabled networks, vehicles do not have stringent constrains in terms of transmission power and can incorporate high computational capabilities to process sensed data into meaningful

information. Also, vehicular mobility on roads is systematically determined by road traffic restriction and are predictable to some extent with vehicle specific information like speed, direction and current position. Despite all these advantages, this concept also has open challenges as vehicles are moving and changing their position constantly which leads to a network topology that varies frequently. Such highly dynamic network with frequent disconnections provides intermittent communication opportunities. In addition, the variable vehicle density on roads makes pure vehicular networks limited to certain geographic areas and to specific times of the day. This mandates frequent use of V2I communication: to that end, understanding the unique vehicular traffic dynamics plays a crucial role in the design and evaluation of networking solution to address the communication properties like latency, loss rate, throughput, offered data traffic load.

1.4 Contribution and outline

In both cases of communication architectures discussed above where vehicles are part of the system, realistic vehicular traffic dynamics are a key requirement. The design and performance evaluation of networking solutions require large-scale testbeds, which is impractical due to cost and complexity concerns. Hence simulation becomes the tool of choice in the validation of such new network architectures and protocols for vehicular environments. Unfortunately, simulative performance evaluation of vehicular networks is often biased by the underlying mobility representation. Early research work in [BSH03, FH08, VBT11b] has proven that the movement of vehicles can dramatically affect the behavior of network protocols, and an incorrect representation of vehicle traffic can lead to misleading conclusions.

In Chapter 2 of this thesis, we first present straightforward solutions involving the collection of real-world vehicular mobility data and highlight their shortcomings. Then, we discuss the need for synthetic vehicular mobility datasets and present the fundamentals of its generation process. We provide a survey of datasets that appeared during the last decade along with a review of the tools that were used in their generation process. Despite many attempts to generate synthetic datasets that meet the realism of real-world vehicular dynamics, our discussion of the properties of an ideal vehicular mobility dataset highlights the problems of the presented traces.

In Chapter 3, we address the need for a large-scale realistic vehicular mobility dataset that closely mimics a typical day vehicular traffic dynamics in an urban area. To that end, we select various open-source state-of-art tools and follow standard approach to generate the vehicular mobility as reviewed in Chapter 2. We discover and solve various types of issues when

integrating these tools. This involves matching the road attributes and macroscopic traffic definitions with the reality, and to use additional tools to achieve realistic traffic assignment and to validate the realism of the dataset. Comparative analysis with other datasets shows that simplistic assumptions on the macroscopic or microscopic dynamics can greatly affect the realism of synthetically generated traffic dynamics. Later, this highly realistic dataset is integrated with the OpenAirInterface emulator platform to represent the mobility of LTE vehicular nodes.

In Chapter 4, we study the impact of a large-scale realistic representation of vehicular mobility on cellular and autonomous network architectures. We first analyze the impact of a pervasive presence of vehicular users on the radio access network. To that end, we obtain the real cellular deployment from the urban area under study and approximate the cellular deployment with a Voronoi tessellation. The macroscopic and microscopic analysis uncovers the mobility dynamics of the vehicular users within the access network. Our results unveil the macroscopic offered load patterns and spatiotemporal flows, as well as the properties of vehicular user processes at network cells. From the view point of autonomous networks, the connectivity analysis stresses that the pure V2V communication seems practical at peak time of the day in a limited geographic area but requires V2I for larger use cases. Our spatiotemporal analysis shows that there exists a relation between mobility patterns, specific areas in a city and time of the day. This let us speculate that this type of patterns could be predicted.

Thus, motivated by the changing mobility patterns in time, in Chapter 5 we propose an online macroscopic prediction approach that is lightweight in the sense that short-time mobility dynamics are used to predict future vehicle movements. The macroscopic nature of this approach profiles the users who follow a similar path. Parameters that affect the prediction performance are discussed and optimal values are presented. The proposed lightweight approach predicts the vehicular traffic flows with 80 to 90% accuracy at crossroads and across cellular cells.

Finally, in Chapter 6, we summarize the main findings of this thesis, the conclusion that can be drawn and the discussion on future possible directions.

2 Synthetic vehicular mobility

Contents

2.1 Introduction 12
2.2 Real-world vehicular mobility datasets 12
2.3 Generation process 14
2.3.1 Road topology database 14
2.3.2 Microscopic traffic flow description 16
2.3.3 Macroscopic road traffic description 18
2.4 Mobility simulators 20
2.4.1 Microscopic traffic simulators 20
2.4.2 Mesoscopic traffic simulators 21
2.4.3 Macroscopic traffic simulators 22
2.4.4 Interactions between simulators 22
2.5 Mobility datasets 24
2.5.1 Perception 24
2.5.2 Small-scale measurements 28
2.5.3 Road traffic imagery 29
2.5.4 Roadside detectors 30
2.5.5 Socio-demographic surveys 31
2.5.6 Discussion 32
2.6 Summary 33

2.1 Introduction

During the last decade, the research on vehicular networks has led to the proposal of countless network solutions specifically designed for vehicular environments. Such solutions encompass all network paradigms with original algorithms and architectures proposed for next-generation roadside cellular access networks, pure vehicular ad hoc networks, and opportunistic disruption-tolerant vehicular networks.

When such vehicular network solutions target vast deployments that comprise tens, hundreds or even thousands of vehicles, simulation is the mean of choice for their performance evaluation. Indeed, large-scale experimental testbeds are too expensive and complex to scale to more than a few communication-enabled vehicles, while comprehensive analytical approaches are often untraceable from a mathematical viewpoint. In that regard, the primary objective becomes to collect and use the vehicular mobility information with the network simulation application. Alternatively, large processing capability in recent computer systems has eased the generation of synthetic mobility datasets of millions of vehicles with realism that mimics the real-world vehicular traffic.

In this Chapter, we present the existing real-world mobility datasets in Sec. 2.2 and discuss their limitations. Later, the process of generating synthetic mobility datasets is discussed in Sec. 2.3, where we also briefly outline the history of synthetic vehicular mobility generation. Sec. 2.4 then provides a taxonomy of the most important tools employed in the process above, i.e., the microscopic-level simulators that are available today for the generation of vehicular mobility datasets. The outcome of the process; i.e., the synthetic mobility datasets that is publicly available for the network simulation, is then presented in Sec. 2.5.

2.2 Real-world vehicular mobility datasets

Many projects involved in tracking and logging vehicle traffic in the real world have first appeared during the first half of the past decade and have since then grown in number and scale. Datasets of this kind with a meaningful scale (i.e., comprising more than a few units) all come from mobile fleets, e.g., bus services or taxi companies. Table. 2.1 highlights the most relevant features of the real-world vehicular tracking efforts.

The first dataset of this kind was recorded in Seattle, WA, USA [JHP+03], and presents the movement of 1200 buses over an area of 5100 km^2 during a period of two weeks. The data was retrieved through the Automatic Vehicle Location (AVL) system, that is commonly deployed in recent public transport systems, with the position of each bus updated every 2

minutes. A similar technique has been more recently employed in [DPPW10] to extract a mobility dataset of 1648 buses traveling in the urban area of Chicago, IL, USA. This second dataset also spans over a long period of time, namely 17 days, and has a time granularity of 20-40 seconds, thus higher than that of the Seattle bus dataset. Both these traces, apart from being limited to buses, are affected by a low time granularity, characteristics that necessarily constrain their use to delay-tolerant networking paradigm.

A similar technique to obtain real-world datasets consists in directly retrieving vehicle positions from GPS receivers. Seminal work in such a direction was conducted in Boston, MA, USA [BGJL06]. The DieselNet testbed consists of 30 buses traveling over a geographical area of 241.40 km^2. The size of the testbed makes it of small interest from a pure mobility viewpoint: in fact, DieselNet is employed to evaluate the communication capabilities of the vehicles, rather than their movement patterns.

A large-scale collection of vehicular GPS recordings is instead the aim of the Shanghai Grid project, where the movement patterns of approximately 4000 taxis and 2000 buses were recorded for several months over an area of 240 km^2 in Shanghai, PRC. From the overall dataset, mobility traces were extracted for taxis [HLL+07] and buses [SLL+08]. In the first case, the dataset captures the movement of 1171 cabs in the Shanghai inner urban area for three months, with a granularity of 1 minute approximately. In the second, the trajectories of 700 buses over a geographical area of 150 km^2 are recorded. Although these datasets involve a significantly increased number of vehicles with respect to the DieselNet testbed, they still remain limited to a minor subset of the overall traffic in the urban area considered, and are affected by coarse time granularity.

Similar features characterize a taxi mobility dataset in Beijing, PRC [ZAKB10], where the positions of 2927 taxis are updated every minute for 24 hours. The scale is instead much larger in the T-Drive dataset [YZXS11], that includes 30000 taxis moving around the same city during three months, although the available refined mobility dataset is limited to 10357 cabs and one week of duration. Finally, the routes of 500 taxis San Francisco, CA, USA, can be fetched from [Cab11], where the data is constantly updated. As for all the other GPS datasets, also these datasets suffer from limited update frequency and penetration rate.

An interesting breakthrough in the collection of real-world vehicular movement data could be provided by the diffusion of recent navigation systems [Tom10, Mei11], that periodically communicate car position information to traffic management centers. Although this may greatly help real-world traffic tracking to scale up, it is easy to foresee that market rules and privacy concerns will hinder the public disclosure of such data, similarly to what happens today with the mobile access network logs of telecom operators. In the light of all these

Real-world dataset	Traffic data	Granularity (s)	Duration (days)	Area (km^2)	Type	Vehicles
Seattle [JHP+03]	AVL/GPS	120	13	5100	bus	1200
Chicago [DPPW10]	AVL/GPS	40	17	606	bus	1648
DieselNet [BGJL06]	GPS	1	60	241.40	bus	30
ShanghaiGrid [HLL+07]	GPS	60	1	102	taxi	1171
Beijing [ZAKB10]	GPS	60	1	750	taxi	2927
T-Drive [YZXS11]	GPS	180	7	750	taxi	10357
San Francisco [Cab11]	GPS	60	30	18000	taxi	500
ShanghaiGrid [SLL+08]	GPS	-	4	150	bus	700
Frankfurt [SIm13]	GPS	-	-	300	cars	420

Table 2.1: Major features of the real-world vehicular mobility datasets.

developments, synthetic vehicular mobility datasets have emerged as the most viable solution to represent node or user dynamics in network simulation.

2.3 Generation process

Like any other system, Fig. 2.1 defines a system that is used to generate the vehicular mobility dataset with three major components mutually influencing each others, i.e., the road topology information composing *Motion Constraints*, and the microscopic flow description and the macroscopic traffic description, composing *Traffic Demands*. Traffic demands provide a description of the desired mobility, whereas motion constraints limit them. In this Section, we introduce such components, outlining their importance to the realism of the resulting movement dataset. Discussing the different components in the order above also gives us the opportunity to review the evolution that the simulation of road traffic for vehicular networking has undergone during the last decade.

2.3.1 Road topology database

Early works on vehicular networks employed simplistic graphs for constrained vehicular mobility, like regular grids [Dav00, BSH03], also referred to as Manhattan scenarios and depicted in Fig. 2.2(a), or simple user-defined layouts [THB+02], as in Fig. 2.2(b). Automatic road topology generation was also explored, by leveraging clustered Voronoi tessellation [HFFB11], which led to layouts such as that portrayed in Fig. 2.2(c). However, all these approaches did not yield a sufficient level of realism, and were soon abandoned.

Use of real-world maps started with the U.S. Census Bureau TIGER database [Bur12], complete topological information on the whole US road network were made freely accessible, hence became especially popular to that end. Notable works employing the TIGER database are

2.3 Generation process

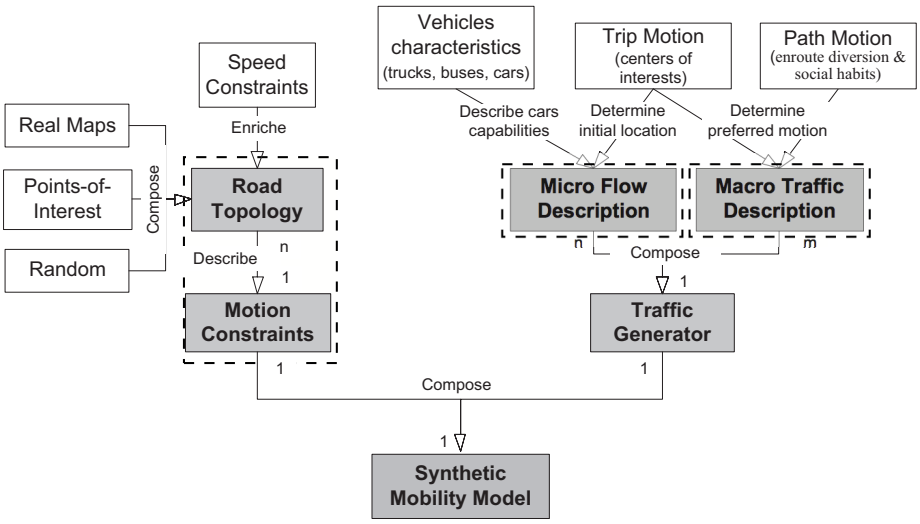

Figure 2.1: Concept map for synthetic vehicular datasets, where the three main blocks (dotted blocks) are enriched with added blocks and features providing an increased accuracy.

those in [SJ04, NTD+06], while an example of TIGER road layout is depicted in Fig. 2.2(d). Another source of road topology information that became quite popular was the Ertico's Geographic Data File (GDF) database [Ert12]. The GDF format was more complete than the TIGER one, but the two databases were complementary, as the Ertico one provided topology information of European cities. The GDF format was employed, e.g., in [HFFB11]. A sample road topology is shown in Fig. 2.2(e).

More recently, the success of the OpenStreetMap (OSM) [Ope12] initiative has granted to the OSM format a de-facto standard status. The OSM initiative applies the open-source concept to road cartography, and thus provides freely exportable maps of cities worldwide, which are contributed and updated by a vast user community. Map information includes roads, railways, buildings, and Points of Interests (PoI) such as parks, commercial centers, leisure centers and commercial activities. The OSM road information is generated and validated by the user community through satellite imagery and GPS datasets, and is commonly regarded as the highest-quality road data publicly available today. Indeed, the accuracy of OSM street layouts, comprising highways, major urban arteries and minor roads, often matches that of proprietary ones such as, e.g., Google Maps or Mappy, especially for large cities. An example of OSM topology is provided in Fig. 2.2(f).

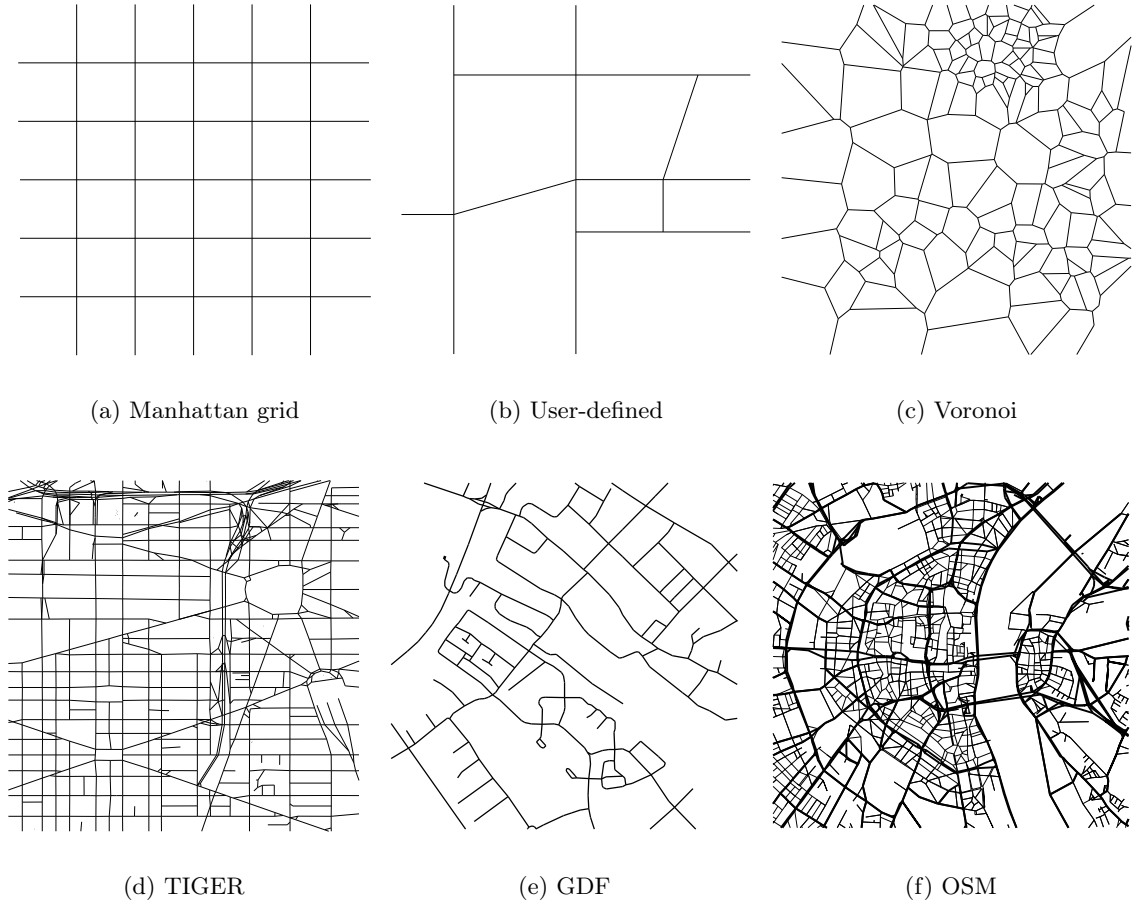

(a) Manhattan grid (b) User-defined (c) Voronoi

(d) TIGER (e) GDF (f) OSM

Figure 2.2: Road topology database examples.

2.3.2 Microscopic traffic flow description

The second key component in the generation of vehicular movement datasets is the microscopic-level mobility model. This model controls the acceleration (possibly deceleration) of an individual car at each time instant, thus determining its speed over time. Diverse microscopic mobility models can have very different levels of detail. Below, we provide an overview of the evolution of microscopic mobility modeling for vehicular network simulation over the last few years. More complete discussion can be found in [M.F09].

Initial works on the simulation of vehicular networks employed *stochastic* microscopic mobility models, that, in their simplest implementation, assigned to each vehicle a random speed in a predetermined speed range. The vehicle then traveled at such a constant speed during its whole trip. Basic stochastic models were employed, e.g., in [Dav00, THB+02, SJ04]. In the latter work, the speed range was made road-dependent, and obtained as $[v_{max} - \delta, v_{max} + \delta]$, where v_{max} was the maximum speed allowed on the current road, and δ a tuneable value.

2.3 Generation process

These models are clearly not representative of a real-world driver behavior, since they completely disregard the fact that the speed of each vehicle depends on the presence of other nearby vehicles. As a result, vehicles can, e.g., travel one over the other and cross an intersection at the same time. A more complex stochastic model was introduced in [BSH03], so as to avoid that vehicles could travel over each other. The proposed model, named *Freeway* model, still employed random speed decision, but forced each driver to maintain a speed lower than that of the preceding vehicle. However, the Freeway model was proven to yield properties very far from the minimum requirements demanded for a realistic microscopic mobility simulation [FH08].

A few microscopic mobility models were proposed by leveraging *traffic stream* properties of road traffic. These models looked at vehicular mobility as a hydrodynamic phenomenon and related the three fundamental variables of velocity v (in km/h), density ρ (vehicles/km), and flow q (in vehicles/h), according to the fundamental equation

$$\frac{\partial \rho}{\partial t} = -\frac{\partial q}{\partial x} = -\frac{\partial (\rho v)}{\partial x}. \tag{2.1}$$

As a result, the speed of a vehicle was determined according to the road traffic intensity over the road segment it was travelling on. Examples of traffic stream models employed in the vehicular networking literature can be found in [SMHW92, RML02]. Although they represent an improvement with respect to stochastic models, traffic stream models yield a limited level of detail, e.g., forcing all the vehicles on a same road segment to travel at the same exact speed. As such, traffic stream models were also proven not to be a solution for the simulation of vehicular networks [FH08].

A significantly higher level of realism is provided by *car-following* models, that describe the behavior of a driver depending on the state of its neighboring vehicles. Most car-following models determine the acceleration of a vehicle in relation to the state of the car in front. A general expression for a car-following model is

$$\frac{dv_i(t)}{dt} = f(v_i(t), v_j(t), x_j(t) - x_i(t)), \tag{2.2}$$

where $v_i(t)$ and $x_i(t)$ are the speed and position of vehicle i at time t. The equation allows to compute the acceleration of i, depending on the absolute and relative speed of both vehicles, as well as their distance $x_j(t) - x_i(t)$. Examples of car-following models employed in the simulations of vehicular networks abound, and include the Intelligent Driver Model (IDM) [THH00] employed, e.g., in [JBL05, HFFB11, BLP08], the Krauss model [KWG97] adopted, e.g., by [Bre01, KML07], or the Gipps model [Gip81] used in [FMR06]. All these

models have been validated through standard tests proposed by the transportation research community.

An alternative to car-following modeling, that is equivalently valid, is represented by *Cellular Automata* models. These models discretize time (by assuming fixed-length time step), space (by fragmenting roads into cells each hosting a single car at most) and speed (by considering a finite state of instantaneous speeds). The movement of cars is then described, at each time step, as a shift of finite states (i.e., vehicles with an associate speed) along a one-dimensional lattice of cells. Different Cellular Automata models implement different rules to realize the shift: a widely adopted model is the Nagel-Schreckenberg one [NS92], used for the simulation of vehicular networks in [FMR06].

All the models discussed to this point deal with *intra-flow* microscopic mobility, i.e., they can be used to determine the movement of vehicles over a single unidirectional lane. However, in most real-world scenarios, *inter-flow* interactions must be modeled as well. The latter occur in highway traffic, when, e.g., one vehicle overtakes the one in front (and thus changes lane, merging into a new flow) or joins the traffic from a ramp. Similarly, when simulating urban traffic, flows constantly merge at intersections or roundabouts. Modeling intra-flows mobility is thus critical to achieve a complete road traffic representation in vehicular network simulation.

Models have thus been proposed that complement the intra-flow descriptions presented above, by adding rules to manage the movement of a vehicles from one flow to another. These inter-flow models are typically built over realistic car-following or cellular automata models, and extend them so as to mimic lane changing and ramp in-flow [TH02, Kra98, NWWS98], as well as intersection management in presence of road signalization (e.g., traffic lights, stop or yield signs) [HFFB11]. A detailed presentation of all these inter-flow models can be found in [M.F09] and references therein for a thorough discussion.

2.3.3 Macroscopic road traffic description

Access to a highly detailed road topology with stop signs, speed limits, traffic lights, movement restrictions along with faithful driver protocols ensuring fluid traffic behavior forms the key tools for synthetic dataset generation. However, an important question remains unanswered: *where should vehicles go?* In other words, the missing component is the macroscopic description of road traffic, that determines when vehicles start their trips and from which location, which route they choose, and where they finally stop.

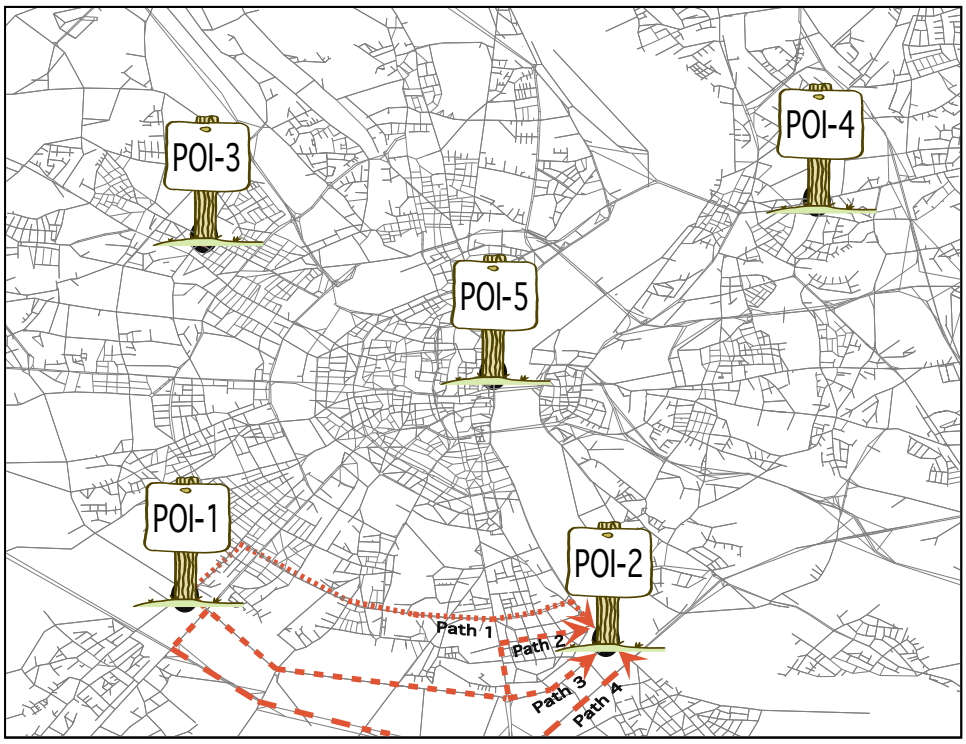

Figure 2.3: Traffic demand and traffic assignment.

The macroscopic description is critical to the realism of the resulting dataset, since it is easy to understand that it determines the volume of traffic in the scenario, as well as its evolution over time. Moreover, the start and stop locations of trips and the routes followed by each vehicle in the real world guide the planing of the road topology, which, as a consequence, is designed to accommodate the real-life traffic flows. Therefore, injecting unrealistic (e.g., random) flows in the road topology risks to lead to large traffic jams in some areas and road under-utilization in others.

In fact, the availability of accurate road topology information, discussed in Sec. 2.3.1, and of detailed intra- and inter-flow mobility models, as from Sec. 2.3.2, makes the description of macroscopic traffic the real challenge on the path towards realistic synthetic datasets of vehicular mobility. In order to faithfully mimic macroscopic traffic flows along highway networks and across metropolitan areas, one must correctly identify the *traffic demand*, i.e., the start time, the origin and the destination of each car trip in the simulated region, represented as a so-called Origin/Destination (O/D) matrix. Fig. 2.3 shows five point-of-interest's (POI's), an O/D matrix contains the travel information between such POI's at a given time. This information is specific to a person hence obtaining such knowledge is only possible through

surveys. Then, an appropriate *traffic assignment* algorithm needs to be run on the O/D matrix, so as to identify the realistic route followed by each driver to reach his/her destination. For example, to reach POI-2 from POI-1, a person can drive through any of the four paths, but to obtain a route that feature less congestion, fuel, and time efficient, one needs a good route assignment algorithm. As traffic capacity threshold of the roads limit huge volume of vehicular traffic choosing the shortest path and is not always a viable option.

Most of the simulative evaluations in the vehicular networking literature neglect the importance of the macroscopic traffic description, and assume random traffic demands, i.e., O/D matrices that contain random locations and times, or fixed probabilistic turns at each intersection. Also, shortest path is the common solution to the traffic assignment problem. Such an approach easily leads to unrealistic flows and thus strongly biased simulation results.

However, as the awareness on the importance of macroscopic modeling grows, solutions start to be proposed and adopted in order to generate datasets that are realistic also from the viewpoint of large-scale traffic flows. In Sec. 2.5, we will thus classify the datasets currently available for the simulation of vehicular networks according to the macroscopic description they employ.

2.4 Mobility simulators

In this section we describe a number of simulators which are capable of integrating the features discussed in Sec. 2.3 to generate synthetic mobility datasets. A larger coverage on vehicular mobility simulators may be found in [Har09, HFC09].

2.4.1 Microscopic traffic simulators

Mostly developed for urban traffic engineering, microscopic traffic simulators such as PARAMICS [PAR12], CORSIM [COR12], VISSIM [VIS12a], TRANSIMS [TRA12], AIMSUN [AIM12], VanetMobiSim [Van12] or SUMO [SUM12], are able to model individual urban traffic at a microscopic level. Yet, not all of these simulators can be used straightaway for vehicular networking, as either synthetic datasets cannot be exported, or open API are not available to interconnect them with a network simulator. The semi-commercial simulators VISSIM and AIMSUN provide support for dataset extractions and interconnection. In the interest of comparability of research results, it is however more beneficial to use readily available Free and Open Source Software simulators, such as SUMO. SUMO is capable of extracting synthetic datasets in the shape of data files, and provides external control through an open API called TraCI.

From a macroscopic road traffic level (topological maps, traffic models), they mostly all support similar features, but it is at the microscopic level that they differ. For instance, TRANSIMS and AIMSUN are based on an extension of the Gipps microscopic flow equations [Gip81], whereas VISSIM is based on a psycho-physical flow model from Wiedemann [Wie74]. SUMO initially supported only the Krauß flow model [KWG97] and its extensions, but has been recently been extended to include the same complex Gibbs model called *Intelligent Driver Model* from Treiber et al. [THH00] as VanetMobiSim.

Microscopic traffic simulators represent the mostly used simulators to simulate vehicular networks, as they are capable of modeling individual mobility at a high level of precision. It is also this feature that made microscopic traffic simulators popular for the design and evaluation of traffic safety applications based on vehicular communications. Scale is their limitations, as modeling individual and strongly correlated mobility at a microscopic level requires a significant computational capability. Nevertheless, simulators such as SUMO or VISSIM are capable of simulating up to 10'000 vehicles faster than or equal to real-time, and SUMO has been the preferred choice for the Open-Source ITS simulation platform iTETRIS [iTe11].

2.4.2 Mesoscopic traffic simulators

One typical approach to simulate traffic at a mesoscopic level is to model road elements as a queue and intersections as nodes. The vehicular dynamics is represented by queue parameters, such as the queue capacity (the maximum bumper-to-bumper vehicles allowed on a road element), and the maximum flow of the queue. The node parameters represent the intersection scheduling policy, such as the intersection type, level of priority and access control to connecting queues. The simulators Matsim [MAT12] and FastTrans [FAS12] are two popular tools in this category. Both simulators support the import of detailed topological maps and detailed added features, and also include efficient traffic models (macroscopic) such as advanced Dijikstra routing and origin-destination matrices. Both are capable of simulating and extracting synthetic datasets of individual cars up to a city-wide scale. FastTrans provides also parallelization on multiprocessors for a higher scalability.

The mesoscopic traffic simulators represent the best trade-off between granularity and scalability, as they can simulate very large scenarios, but still providing the detailed mobility values of each vehicle. For instance, mesoscopic simulators are a preferred choice to simulate large scale impacts of vehicular communication to improve traffic efficiency. Mesoscopic simulators are yet not recommended for the evaluation of vehicular networking for traffic safety applications, as they do not provide a sufficient granularity, in particular between vehicles located in the same mesoscopic queue.

2.4.3 Macroscopic traffic simulators

Simulators in this category concentrates on quantities of macroscopic meanings such as flow, speed, or density instead of a specific vehicle as seen in microscopic simulators. Inspired by fluid theory, macroscopic flows have the advantage of a reduced computational complexity compared to the microscopic ones but they can still realistically model macroscopic quantities. On the other hand, the synthetic datasets obtained by macroscopic simulators cannot directly be used for vehicular networking, as they do not provide detailed mobility values of single entities, but instead aggregated measures such as density and speed. Macroscopic simulators are also unable to model mobility changes or personalized navigation of single vehicles that is typically found in ITS applications for traffic efficiency.

The most well known simulator capable of simulating vehicular traffic at a macroscopic scale is VISUM from PTV [VIS12b]. As other simulator, it is capable of importing precise topological maps, and can configure the traffic light policies at a very precise level. It is typically used by urban planners and traffic departments at the city and regional level to evaluate traffic congestions. Also, macroscopic synthetic datasets are often used to calibrate the microscopic traffic simulators, such as SUMO. Within the iTETRIS project [iTe11], for instance, traffic in the city of Bologna has been simulated using SUMO, but during the calibration step, synthetic traffic flows extracted from VISUM have been used.

2.4.4 Interactions between simulators

Numerous approaches and models have been embedded into various types of simulators which make them unique in its kind. Hence taking advantage of the best features from various simulators would probably help in obtaining close-to-reality synthetic traffic datasets. Traffic calibration first requires multiple sources of data that comes from external sources. Second, vehicular networking needs to be able to use the mobility information from the simulators. For the sake of clarity in this Chapter, we will differentiate between the use of multiple traffic simulators to calibrate one or multiple simulators to obtain synthetic datasets, and the use of multiple types of simulators for vehicular networking. But conceptually speaking, both categories fall into the same domain relating the interactions between simulators.

Calibrating simulators

The first step to design and obtain synthetic traffic datasets is to select the traffic simulator fitting best. Yet, as any tool, a simulator needs to be fed with parameters, such as topological map, traffic demand, traffic lights policy or driver models. Such data may either be obtained

from statistical behaviors, or extracted from dedicated simulators. Considering the case of the calibration of the city of Bologna, as conducted during the iTETRIS project [iTe11], the selected traffic simulator was SUMO. Yet, multiple various other simulators have been used to be able to calibrate the synthetic datasets. Conceptually speaking, calibrating a simulator may be represented as in Figure 2.4. First, the macroscopic simulator VISUM has been configured with the topological map of the city of Bologna and from vehicular flow interactions measured at intersections. The output of VISUM took the shape of traffic volumes at each road segment that have later been imported by SUMO to calibrate its traffic demand. Second, a traffic light control simulator has been used to be able to provide close-to-reality traffic light policy in the city of Bologna. These policy has then been integrated into SUMO. The output has been a calibrated synthetic dataset of the typical traffic in the city of Bologna for 1 hour.

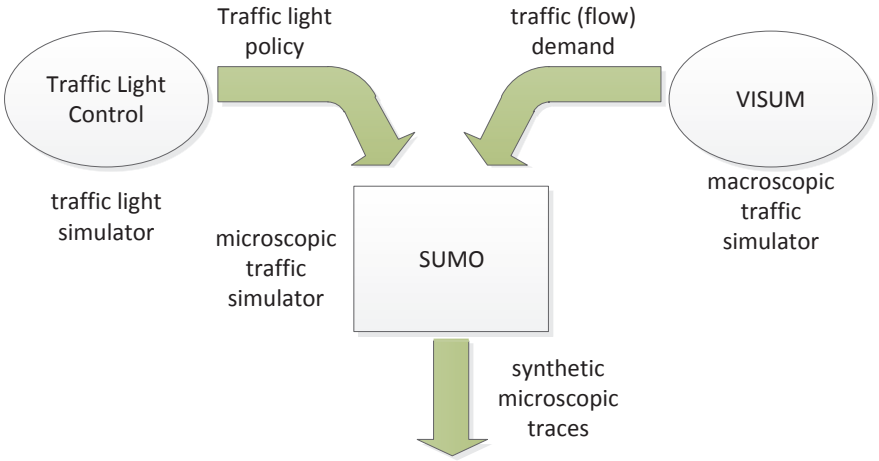

Figure 2.4: The microscopic simulator SUMO calibrated by external simulators.

Interaction between simulators for vehicular networking

We previously described the different methods required to obtain close-to-reality synthetic vehicular datasets. Yet, in order to be used by the vehicular networking community, these datasets need to be made available to network simulators. The description of network simulators available to vehicular research is out of the scope of this Chapter, but we still may cite ns-3 [NS312] and OMNET++ [OMN12] as two popular simulators that are often used in research, as cited in [JDS12].

The interaction between traffic and network simulators may be categorized into two groups as illustrated in Fig. 2.5: *Isolated* and *Federated*. In the former case, synthetic datasets are extracted from traffic simulators in the shape of mobility files that may later be integrated into the network simulator using a dedicated parser. In the later case, a bi-directional interaction is

created between two (or more) simulators to be able to exchange mobility data or to influence the control of one or the other. This approach is more complex, as an interface has to be developed to synchronize and exchange data.

The *Isolated* case has initially been the preferred choice for network design based on small scale random mobility, but nowadays, network simulators are able to integrate large-scale realistic mobility data files from various sources. This approach also remains a favorite choice when mobility does not need to be influenced by vehicular communication or networking. Considering applications for Intelligent Transportation Systems (ITS), mobility shall be influenced, either to avoid accident, or to reduce congestion. So, the *Federated* case is mostly used in that case. Several integrated tools have been created and made available to the community. The iTETRIS platform interlinks the simulator SUMO and ns-3 [NS312], the VEINS platform [VEI12] interlinks SUMO and OMNET++ [OMN12], and VSIMRTI [VSI12] provides APIs to interlink jist/swans [SWA12] and VISSIM. iTETIS and VEINS are both open-source, but VsimRTI only provides open APIs without releasing the code of its core.

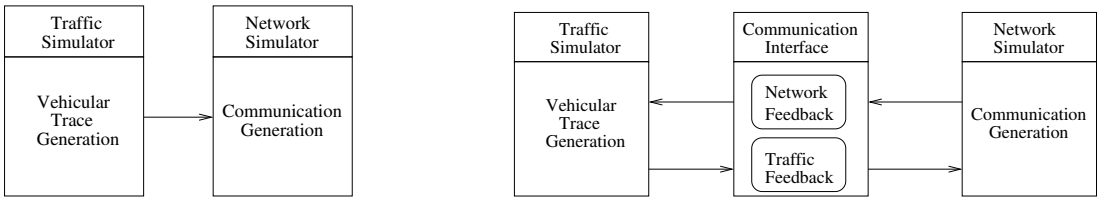

(a) Isolated Case

(b) Federated Case

Figure 2.5: Interactions between network and traffic simulators

2.5 Mobility datasets

In this Section, we overview the body of work on vehicular mobility datasets 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. The relevant features of the different mobility datasets are summarized in Tab. 2.2.

2.5.1 Perception

A simple way to fill the O/D matrix and determine the traffic assignment is to employ one's perception of the road traffic. The authors of [BG09] leveraged their knowledge of the traffic in Turin, Italy, so as to determine the macroscopic demand injected in different areas of the

Type / Name	Simulator	Duration (h)	Area (km^2)	Vehicles	Granularity (s)	Roads
Perception						
Turin [BG09]	SUMO	1	20	1500	1	major/minor
Porto [CaDFJ08]	DIVERT	0.3	62	5000	1	major/minor
Zurich [BLP08]	GMSF	0.5	9	420	1	major/minor
Measurements						
Karlsruhe [FCCP12]	VanetMobiSim	0.7	21	2000	1	major/minor
Berlin [LHT+03]	Videlio	14	21.56	955	1	major/minor
Photography						
Porto [FCFO09]	-	-	41.3	10566	snapshot	all
Detectors						
Berkeley [BK09]	-	72	-	-	60	single freeway
Toronto [BK09]	-	24	-	-	20	single freeway
Bologna [iTe11]	SUMO	1	20.6	10333	1	major/minor
Luxembourg [PDP11]	SUMO	12	1700	150000	1	major
Survey						
Portland [WK99]	TRANSIMS	0.25	21	16529	1	major/minor
Canton of Zurich [RCV+03]	MMTS	6+6	65000	260000	1	major

Table 2.2: Major features of publicly available synthetic vehicular mobility datasets

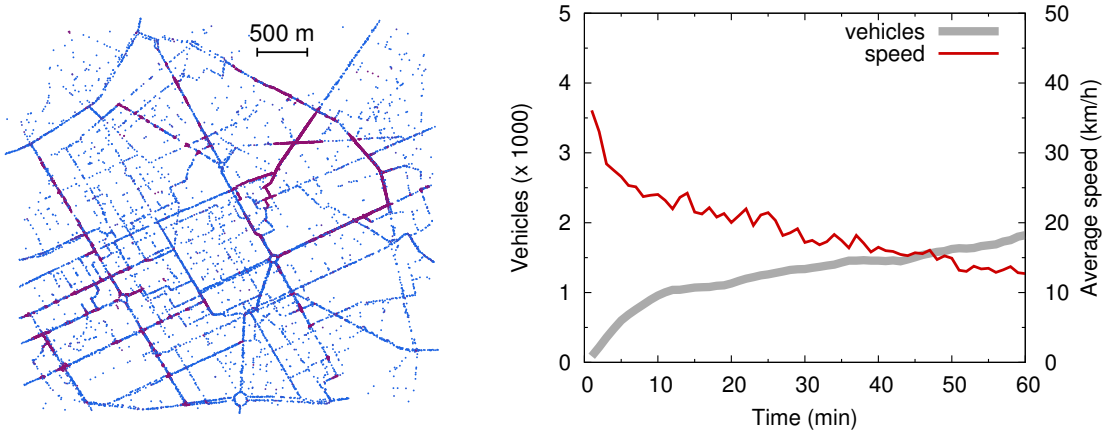

(a) Traffic snapshot

(b) Traffic volume and speed

Figure 2.6: Turin scenario: downtown, sparse traffic. This figure is best viewed in colors.

city. The datasets were generated using OSM as the road topology database and SUMO as the microscopic traffic simulator. Fig. 2.6 and Fig. 2.7 refer to the *Turin downtown* and *Turin outskirts* scenarios, respectively. In both figures, the left plots portrays a snapshot of the vehicular mobility, so as to provide an intuitive glimpse of the mobility reproduced by each dataset. Also, colors are used to distinguish fluid traffic conditions from traffic congestion: light blue denotes high average speeds, whereas dark violet indicates speeds close to zero. We can observe heavy localized traffic jams in Fig. 2.6(a): given the sparse traffic conditions assumed in this scenario, congestion is an artifact of the incorrect traffic demand modeling rather than a representation of a real world phenomenon. The snapshot in Fig. 2.7(a) depicts instead a much simpler road topology, where traffic is fluid along the few major roads and absent elsewhere. Congestion is only recorded around the main roundabout in the northern region of the scenario. The right plots of Fig. 2.6 and Fig. 2.7 show instead the time evolution of the traffic volume and average speed, for the same two scenarios. We can observe that, notwithstanding the limited duration of the datasets, none of them is in stationary conditions. In both cases, the number of vehicles in the region tends to grow, while speeds decrease. The significant difference between the traffic conditions at the beginning and at the end of the datasets is another evidence of the difficulty of calibrating the traffic demand based on pure perception. As an example, in Fig. 2.6(b), the average speed starts at 30 km/h (that, considering the presence of traffic lights, denotes a completely empty road network) and progressively drops to 10 km/h (that is instead a symptom of major congestion), in less than one hour. Such a phenomenon would be hardly observed in the real world, in presence of the sparse traffic assumed in the dataset.

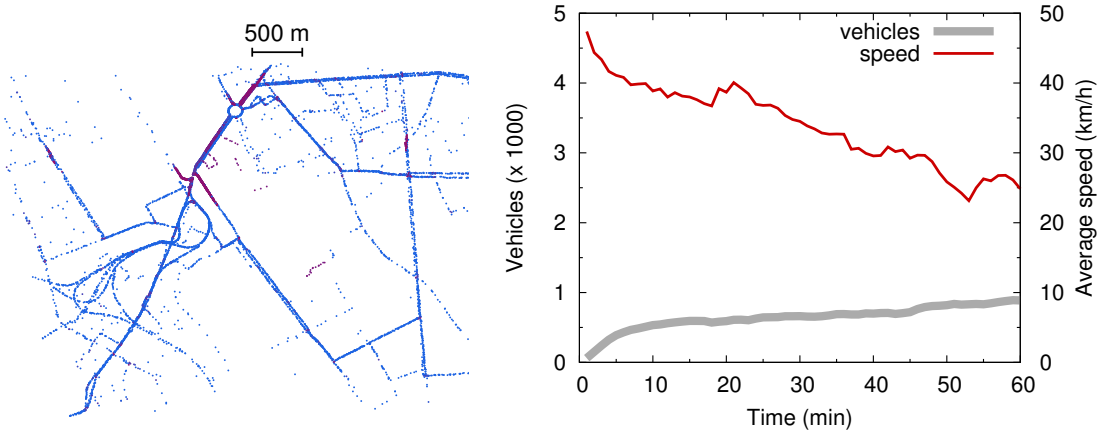

(a) Traffic snapshot

(b) Traffic volume and speed

Figure 2.7: Turin scenario: city outskirts, dense traffic. This figure is best viewed in colors.

We remark that a similar macroscopic approach is also adopted in [CaDFJ08], although a different microscopic simulator framework, named DIVERT, is employed. The city concerned is in this case Porto, in Portugal.

The technique used to generate mobility datasets of the region of Zurich, Switzerland, in [BLP08], is instead slightly different. Apart from the microscopic simulation environment adopted, that is this time Generic Mobility Simulation Framework (GMSF), and the road topology information, retrieved from the Swiss Geographic Information System (GIS), the macroscopic traffic flows are injected assuming that larger roads attract more traffic. The features of the resulting mobility are reported in Fig. 2.8, for one of the scenarios in [BLP08], namely that representing the city center (the others refer instead to suburban and rural areas), which we name it as *Zurich downtown*. From the traffic snapshot, in Fig. 2.8(a), we can observe that, in the small region compassed by the simulation (less than 10 km^2), the traffic tends to be consistently slow over the whole street layout. Such a uniformity is also present in time, as the number of vehicles in the region and the average speed are constant over the 20 minutes of duration of the dataset. Overall, the Zurich city downtown dataset exhibits speeds exceedingly low with respect to the traffic volume considered. More importantly, the mobility description appears significantly limited in terms of size and temporal duration.

In fact, a common weakness of datasets built on simplistic macroscopic traffic descriptions is that they show limited extension, in both space and time. As a matter of fact, they only cover a few tens of km^2, and have durations ranging between 20 minutes and one hour. This limitation is a by-product of the macroscopic description: the perception approach can be

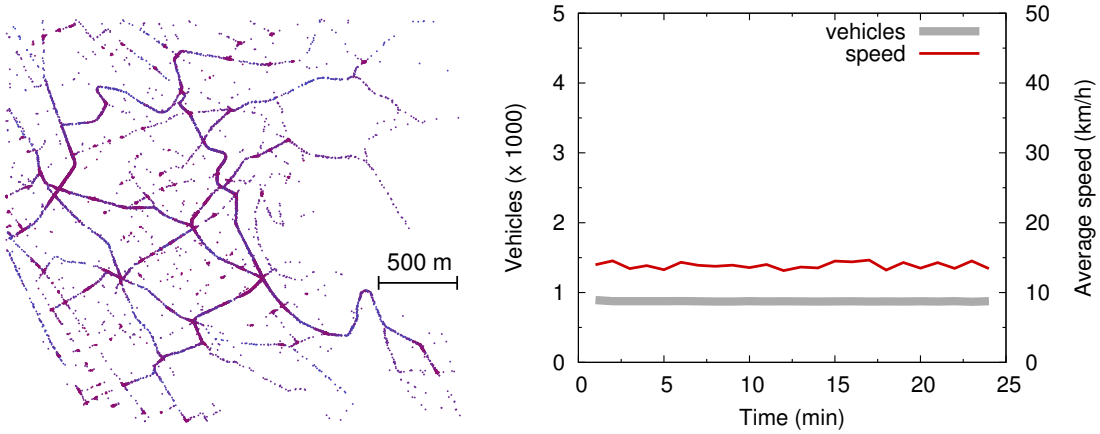

(a) Traffic snapshot

(b) Traffic volume and speed

Figure 2.8: Zurich scenario: downtown. This figure is best viewed in colors.

applied on small areas for short time periods. However, the lack of real-world data behind the traffic demand and assignment leads to significant problems when large regions or long time periods are taken into account. Manually assigning large-scale flows so that the road topology can accommodate them is a prohibitive task. Similarly, the introduction of realistic variations of road traffic over time is hard to accomplish, if no realistic data is available to support it.

2.5.2 Small-scale measurements

Several works have exploited small-scale measurements, typically conducted by the authors themselves, so as to provide some simple validation to the authors' perception, or to complement the latter. This is the case of the mobility dataset presented in [LHT+03], that reproduces the vehicular traffic in the downtown Berlin, Germany. The dataset, generated with a dedicated microscopic simulator named Videlio, covers a small area of around 20 km^2 for a quite long period of 14 hours. However, the low number of vehicles (less than 1000) present in the dataset over such a long time interval raises question on the realism of the macroscopic traffic demand employed.

Small-scale measurements are also at the base of the macroscopic modeling in the dataset of downtown Karlsruhe, Germany, employed in [FCCP12]. The dataset, generated by feeding a OSM map to the VanetMobiSim microscopic mobility simulator, mimics the movement of 2000 vehicles for approximately 40 minutes. The plots in Fig. 2.9 show that, similarly to what observed for the Zurich datasets before, the number of vehicles is quite stable, while the

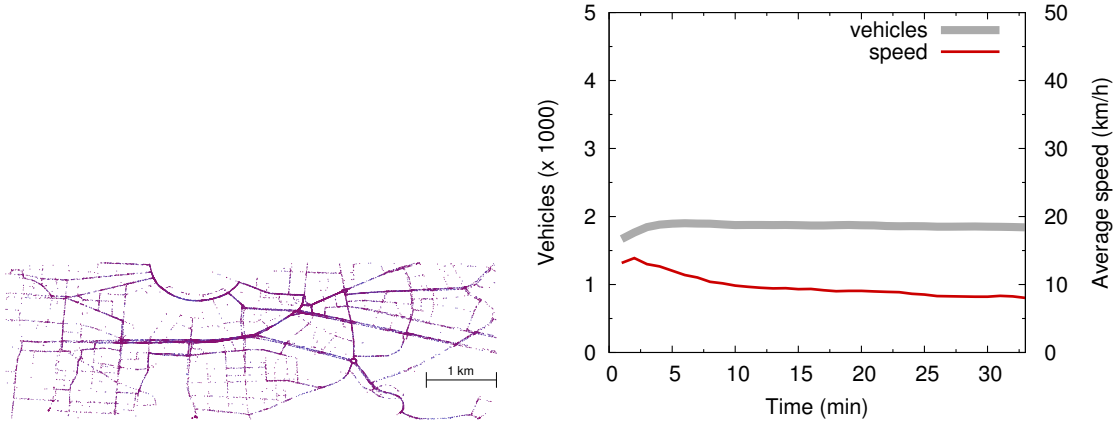

(a) Traffic snapshot

(b) Traffic volume and speed

Figure 2.9: Karlsruhe scenario: downtown. This figure is best viewed in colors.

vehicular mobility is characterized by a generalized low speed. In fact, the average velocity recorded in the dataset shows a slow but constantly decreasing trend with a steady traffic volume, a symptom of potential problems when extending the simulation duration.

Overall, approaches based on small-scale measurement represent an appreciable effort towards the simulation of more realistic traffic demands. However, it is clear that they lack the statistical rigor needed for a realistic representation of the macroscopic traffic distribution, and, as such, can hardly capture the complexity of traffic flows in urban areas or their evolution over long time periods. For these reasons, also these datasets only cover modest geographical surfaces, constrained to a few tens of square kilometers, or have limited time duration, in the order of tens of minutes.

2.5.3 Road traffic imagery

An original solution to the derivation of the macrocopic traffic information was adopted in [FCFO09], where the authors leveraged stereoscopic aerial photography in order to capture the vehicle distribution in the city of Porto, Portugal. A private aircraft was flown over the city for two hours in the early afternoon of a weekday, and photographs were shot from the plane every 5 seconds. The flight followed a parallel row pattern so as to cover the whole geographical area of 41.3 km^2 corresponding to the surface of Porto. By studying the aerial imagery, the authors were able to reconstruct a single snapshot of the positions of 10566 vehicles in the urban area. The time error in the snapshot, due to the fact that pictures of different locations were taken at different moments, is of 23 minutes between two cars, on average. Although this appears as an interesting way to derive static macroscopic data, its

applicability to the generation of mobility datasets is not clear, due to the complexity and cost of running a large-scale aerial photography campaign for a significant period of time.

Another interesting attempt at using imagery to estimate the macroscopic behavior of car traffic has been recently presented in [THKH11]. There, the authors exploit the pervasiveness of road surveillance cameras to infer traffic densities in two large urban areas, namely London and Sydney. Although promising, this approach has not yet been used for the generation of traffic demands in mobility datasets.

2.5.4 Roadside detectors

Induction loops, infrared counters and roadside sensors represent the traditional way to measure vehicular traffic flows in both freeways and urban road networks. The authors of [BK09] use two sets of empirical data obtained from dual-loop and metal detectors from sections of the I-80 Freeway in Berkeley, CA, USA, and of the Gardiner Expressway in Toronto, Canada. The detector information covers a span of 24 hours and allows to determine the per-lane inter-vehicle arrival time and spacing. Although the data could be fed to a microscopic simulator to derive the position of individual vehicles over time, its validity is limited to highway environments.

Roadside detectors have been employed in an urban environment within the iTetris project [iTe11]. Vehicular mobility datasets of several areas of the city of Bologna were generated accounting for macroscopic traffic data acquired through 636 induction loops spread over the road network, and complemented by user surveys on usual commuting trips. The main dataset covers 20.6 km^2 in the city center for a period of one hour, featuring the movement of 10333 vehicles. Thanks to the real-world nature of the macroscopic data they are built upon, these datasets reach an unprecedented level of realism. Unfortunately, they do not cover large surfaces nor long time periods. Moreover, public accessibility to the mobility datasets not yet granted by the project consortium.

A similar approach has been taken in [PDP11], where the authors calibrate the microscopic mobility simulation of the city of Luxembourg through traffic flow information gathered by the local Ministry of Transport. As such real-world data only covers major traffic arteries, it is complemented by driver routes inferred from the different nature of geographical zones in the area under study, and used to define traffic flows on medium- and small-sized roads. The resulting mobility dataset covers a very large area of 1700 km^2 and features 150000 car trips. Although a very interesting dataset, the Luxembourg dataset focuses on highway and major roads traffic, is limited to the morning period, and only accounts for inbound flows, i.e., traffic moving towards the city center.

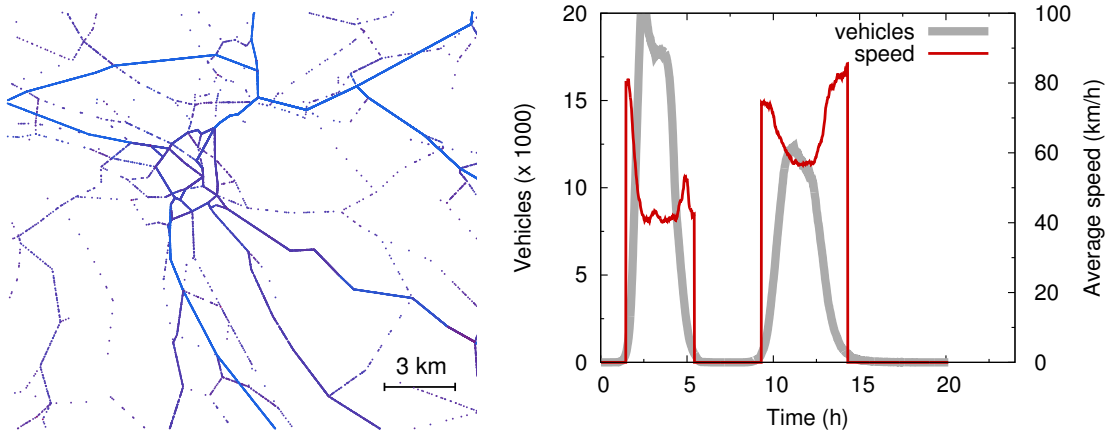

(a) Traffic snapshot

(b) Traffic volume and speed

Figure 2.10: Zurich scenario: region (city and surroundings). This figure is best viewed in colors.

2.5.5 Socio-demographic surveys

Socio-demographic surveys represent a significant source of information for the derivation of vehicular traffic data. The seminal work in [WK99] presents a synthetic mobility dataset whose macroscopic model is derived by knowledge of drivers' activity in downtown Portland, OR, USA. The resulting mobility dataset is acknowledged to be very realistic, but only covers 15 minutes of car traffic in an area of 21 km^2, for a total of 16529 simulated cars. Moreover, the dataset is not publicly available.

The largest vehicular mobility dataset generated to date is the one reproducing the road traffic in the whole Canton of Zurich, a 65000 km^2 region of Switzerland [RCV$^+$03]. The mobility dataset is said to cover 24 hours, through traffic demand data obtained from the Swiss Regional Planning Authority and complemented using the 1994 Swiss National Travel Survey. The resulting dataset, as well as subsets of the same, are widely employed in the vehicular networking literature. However, the size of the road topology forced the authors to limit the detail of the microscopic-level simulation. Thus, they resorted to a queue-based Multi-agent Microscopic Traffic Simulator (MMTS), significantly less accurate than standard fine-grained vehicular mobility simulators based on car-following models. Moreover, once more due to scalability reasons, the road topology was pruned down to major traffic arteries, as observable in Fig. 2.10(a), that portrays a snapshot of the vehicular mobility in the most significative 400 km^2 portion of the simulated region, we name this dataset as *Zurich region*. The selected area represents the city of Zurich and its surroundings: the low topological

detail is striking, especially when comparing the figure with Fig. 2.8(a), that represents the downtown only, but at a much higher detail level. The light blue color that dominates the plot also indicates a high average speed: this is normal, since only major roads, i.e., highways and freeways, are represented. However, the lack of minor urban and rural roads clearly affects the completeness of the mobility description provided by the dataset.

Elevate average speeds are also observed in Fig. 2.10(b), where, however, the more interesting aspect is the limited detail of the traffic demand. As a matter of fact, the plot evidences how the dataset only considers vehicular mobility during the morning and afternoon traffic peaks, while the off-peak hours are neglected, with the traffic volume dropping to zero. The Canton of Zurich dataset is therefore a clear example of the tradeoff between complexity and scale in the generation of synthetic vehicular mobility datasets: in order to be able to reproduce road traffic on a very large-scale, one has to accept low road topology detail, low microscopic precision, and approximate macroscopic description.

2.5.6 Discussion

Overall, an ideal vehicular mobility dataset for network simulation should feature all of the following:

- represent the completeness of road traffic;
- have a high time granularity, the position of each vehicle being traced with a order-of-second precision at least;
- compass very large regions (i.e., whole urban areas), with a faithful description of the road layout and signalization;
- provide a realistic representation of the microscopic behavior of individual drivers in everyday traffic, accounting for their interactions with other drivers and their behavior with respect to road regulations;
- be realistic also from a macroscopic point of view, by faithfully mimicking the movement of large-scale traffic flows across a metropolitan area, over long time periods.

The first two requirements rule out datasets obtained though real-world tracking, as they are limited to subsets of the vehicles, i.e., buses or taxis, and exhibit reduced temporal detail, i.e., order-of-minute position information updates. These datasets are today mainly employed for the performance evaluation delay-tolerant or opportunistic data exchanges, and remain interesting for applications that fit such transfer paradigms.

In the context of synthetic dataset generation, also the third constraint can be easily met, thanks to the availability of accurate real-world road map services. Similarly, recent vehic-

ular mobility simulators allow to satisfy the fourth requirement, as they implement realistic microscopic vehicular mobility models, borrowed from transportation theory and validated through real-world observations. Today's challenge lies indeed in the last aspect, i.e., attaining macroscopic-level realism in the simulation of road traffic. In transportation theory, the problem is separated into two phases. First, one has to identify the *traffic demand*, i.e., the start time, the origin and the destination of each car trip in the simulated region, which are stored in a O/D matrix. Then, an appropriate *traffic assignment* model needs to be run on the O/D matrix, so as to identify the realistic route followed by each driver to reach his/her destination.

The synthetic datasets presented in this Chapter address the issue of macroscopic realism by using different data sources, as clearly outlined by the classification we proposed. Some sources are hardly credible by their own nature (e.g., perception or small-scale measurements) or not scalable due to their complexity (e.g., aerial photography). In the end, and quite unsurprisingly, the most reliable and scalable sources are those commonly employed by transportation engineering (e.g., road traffic detectors and population surveys).

However, none of the datasets making use of realistic macroscopic traffic data sources covers large urban regions for long periods of time, and provides at the same time high microscopic-level accuracy. The Berkeley and Toronto datasets [BK09] are limited to one freeway segment. The Bologna [iTe11] and Portland [WK99] datasets focus on small urban areas of a few tens of km^2. The Luxembourg [PDP11] and Canton of Zurich [RCV+03] datasets are those closer to compassing all of the desirable features we outlined above. However, they mainly model traffic on major roads, and are limited to in-bound flows (Luxembourg) or to traffic peak hours (Canton of Zurich). Hence there is a need for a realistic vehicular mobility dataset which satisfies the points mention earlier in the discussion.

2.6 Summary

In this Chapter, we presented the discussion regarding the basic approaches involved in generation of synthetic vehicular mobility dataset. Various state-of-art tools were discussed along with their history of adoption to date. A brief survey on the available mobility datasets categorized by the nature of the macroscopic information they adopt were presented along with their properties. Despite many attempts to generate a realistic vehicular mobility dataset, the road to the generation of ultimate synthetic vehicular mobility datasets is still long. More precisely, aspects that are still to be addressed concern at a time the scalability of microscopic simulation and the increased detail of the same. Moreover, the realism of macroscopic traffic

description remains an open problem, because realistic road traffic demand information on vast regions is hard to retrieve. To that end, the growing trend of government bodies and local administrations making urban-related data publicly available can provide a significative breakthrough. Finally, we conclude with a discussion on the requirements of an ideal dataset that can serve best either for validation of new network architecture or for protocol evaluation.

A large-scale realistic vehicular mobility scenario: Koln, Germany

3

Contents

3.1 Introduction 36
3.2 Dataset generation process 36
3.2.1 Road topology 36
3.2.2 Microscopic vehicular mobility 37
3.2.3 Traffic demand 37
3.2.4 Traffic assignment 37
3.2.5 Simulation 38
3.3 Repairing the dataset 39
3.3.1 Over-comprehensive and bursty traffic demand 40
3.3.2 Inconsistent road information 42
3.3.3 Flawed road topology conversion 44
3.3.4 Simplistic default traffic assignment 46
3.4 Koln vehicular mobility dataset 47
3.5 Integrating the Koln dataset into OpenAirInterface 53
3.6 Summary 55

3.1 Introduction

The survey of the synthetic mobility dataset generation tools, procedures and existing mobility datasets in 2 showed that there exist a need for a large-scale, highly detailed realistic vehicular mobility description. Specifically, the discussion in Sec. 2.5 outlines the most significant features that any mobility datasets need to be encoded with for more general use cases.

In this Chapter, we introduce a new vehicular mobility dataset of the city of Koln, Germany. To that end, we detail the features of state-of-art tools in Sec. 3.2. Challenges faced during the integration of tools and ways adopted to address it are presented in Sec. 3.3. The resulting realistic large-scale vehicular mobility dataset is discussed in Sec. 3.4 with reference to the mobility datasets presented in Sec. 2.5. Sec. 3.5 highlights the integration of Koln dataset with a LTE emulator: OpenAirInterface. This demonstrates the validity of the synthetic mobility dataset generation procedures discussed in Chapter 2 and uncovers the potential challenges dealing with the integration of various state-of-art data sources and simulation tools.

3.2 Dataset generation process

In this section we employ the tools which are essential for the generation of vehicular mobility dataset as mentioned in previous Chapter. We also discuss the additional procedures which were needed to achieve the greater level of realism.

3.2.1 Road topology

The street layout of the Koln urban area is obtained from the OSM database [Ope12]. In addition to the discussion presented in Sec. 2.3.1, OSM also encourages the use of geographical data in various formats and provides various tools to correct any imperfection encountered in the map data. With its vast user community, OSM always out stands other map providers in availability, ease of use, supporting tools and amount of information.

We employ the Osmosis tool [Osm12] to filter the OSM data and extract the road topology information for an area of approximately 400 km^2 around the urban agglomeration of Koln, thus including almost 4500 km of roads. We then resort to the Java OSM Editor (JOSM) [JOS12] to repair the OSM data file and make it compatible with the microscopic mobility simulator, as detailed in Sec. 3.3.

3.2.2 Microscopic vehicular mobility

The microscopic mobility of vehicles is simulated with the Simulation of Urban Mobility (SUMO) software [SUM12]. SUMO is an open-source, space-continuous, discrete-time traffic simulator developed by the German Aerospace Center (DLR), capable of accurately modeling the behavior of individual drivers, accounting for car-to-car and car-to-road signalization interactions. More precisely, SUMO can import road maps and information on traffic lights, roundabouts, stop and yield signs from multiple formats, including OSM. The microscopic mobility models implemented by SUMO are Krauss' car-following model [KWG97] and Krajzewicz's lane-changing model [Kra09], that respectively regulate each driver's acceleration and overtaking decisions, by taking into account a number of factors, such as the distance to the leading vehicle, the traveling speed, and the acceleration and deceleration profiles. These models have been long validated by the transportation research community, a fact that, jointly with the high scalability of the simulator, makes of SUMO the most complete and reliable among today's open-source microscopic vehicular mobility generators. The version we employed for the dataset generation is 12.3.

3.2.3 Traffic demand

The traffic demand information on the macroscopic traffic flows across the Koln urban area are derived through the Travel and Activity PAtterns Simulation (TAPAS) methodology [VW06]. This technique generates the O/D matrix by exploiting information on (i) the population, i.e., home locations and socio-demographic characteristics, (ii) the points of interests in the urban area, i.e., places where working and free-time activities take place, and (iii) the time use patterns, i.e., habits of the local residents in organizing their daily schedule [HW04]. Within the context of the TAPASCologne project, the aforementioned TAPAS methodology is applied on real-world data collected in the Koln region by the German Federal Statistical Office, including 30700 daily activity reports from more than 7000 households [GR02, EB96]. The resulting O/D matrix faithfully mimics the daily movements of inhabitants of the area for a period of 24 hours, for a total of 1.2 million individual trips. The TAPASCologne O/D matrix is, to the best of our knowledge, the only realistic traffic demand dataset of a large urban region available to date.

3.2.4 Traffic assignment

The actual assignment of the vehicular traffic flows described by the TAPASCologne O/D matrix over the road topology is performed by means of Gawron's algorithm [Gaw98]. This

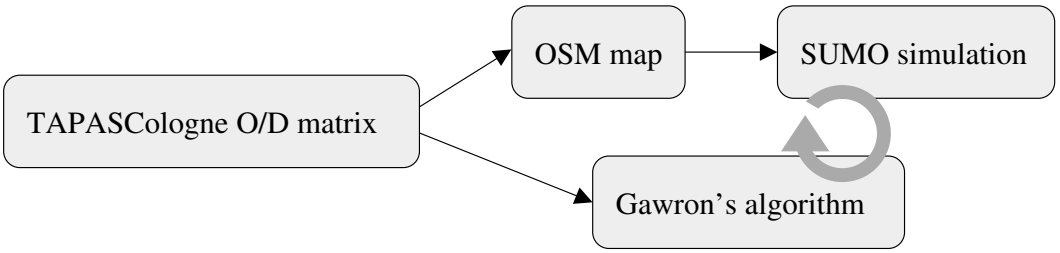

Figure 3.1: Simulation workflow.

traffic assignment technique computes the fastest route for each vehicle, and then assigns to each road segment a cost reflecting the intensity of traffic over it. By iteratively moving part of the traffic to alternate, less congested paths, and recomputing the road costs, the scheme finally achieves a so-called *user equilibrium*. Additionally, since the intensity of the traffic demand varies over a day, the traffic assignment model must also be able to adapt to the time-varying traffic conditions. Indeed, Gawron's algorithm satisfies such a requirement, thus attaining a so-called *dynamic user equilibrium*. Gawron's is one the most popular traffic assignment techniques developed within the transportation research community, and allows to reach a road capacity utilization close to reality and significantly higher than that obtained with, e.g., a standard weighted Dijkstra algorithm.

3.2.5 Simulation

The individual components presented above are combined as depicted in Fig. 3.1 in order to generate the vehicular mobility dataset. First, the information contained in the TAPAS-Cologne O/D matrix are 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 TAPASCologne 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 traffic assignment is generated that allows to sustain the whole volume of the traffic demand.

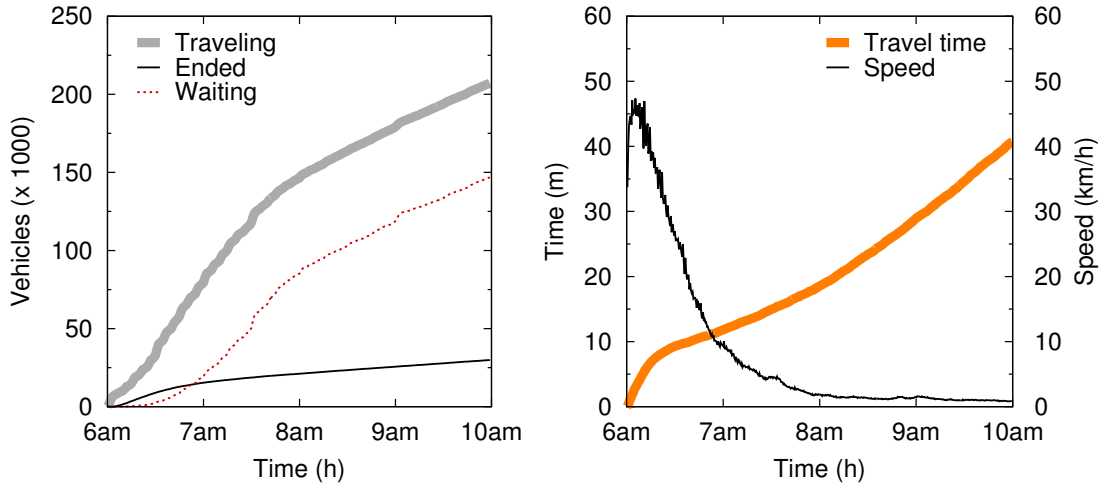

(a) Simulated vehicles status

(b) Mean travel time and speed

Figure 3.2: Original TAPASCologne dataset: Traffic features over time.

3.3 Repairing the dataset

Making all of the previous components work together is not a trivial task. Indeed, when simply running the SUMO simulation with the data sources made available by OSM and TAPASCologne, the result is plain unusable. In Fig. 3.2(a), we plot the time evolution of the number of vehicles that (i) are traveling on the road topology, (ii) have successfully ended their trip by reaching the their destination, (iii) are waiting to enter the road topology, which they cannot presently do due to an excessive congestion of the road segment they are supposed to start their trip from. This third condition is a simulation artifact, identifying situations where the road topology cannot accommodate all the traffic demand in the O/D matrix, and, as such, is an undesirable effect. Clearly, the three sets of vehicles above are disjoint.

From the plot, we can note how the number of traveling vehicles present in the simulation rapidly grows up to exceed a hundred thousands units, a figure completely unrealistic for a city the size of Koln. Additionally, such a number does not tend to decrease as one could expect once the morning traffic peak is exhausted; instead, it keeps growing indefinitely. It is also possible to observe that the number of vehicles that end their trip grows very slowly over time: in fact, from the values portrayed in the figure, only a very small fraction of the cars that are present on the road topology can reach their destination. Finally, the number of vehicles that are waiting to enter the road topology, which we would like to stay as close as possible to zero, grows to hundreds of thousands of units.

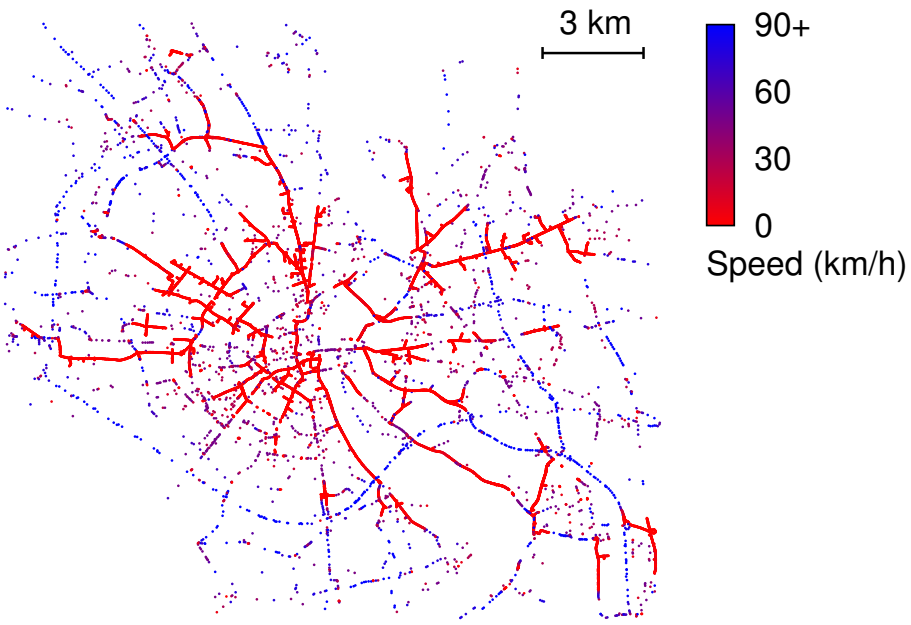

Figure 3.3: Original TAPASCologne dataset. Snapshot of the traffic status at 7:00 am, in a 400 km^2 region centered on the city of Koln. This figure is best viewed in colors.

The mean travel time, in Fig. 3.2(b), also shows a quite unrealistic behavior, as more than half an hour is required, on average, for a driver to reach its destination at 10:00 am, when the traffic should be sparse. Similarly, the average speed of vehicles, in the same plot, tends to zero as the time elapses.

These results are clear symptoms of how the road topology cannot sustain the volume of cars injected according to the traffic demand model. Indeed, when looking at a snapshot of the car traffic in the region, it is evident how the simulation quickly reduces to a huge traffic jam. As an example, Fig. 3.3 depicts the map of the road traffic at 7:00 am: the road topology is mostly covered by bright red dots, representing cars stuck in heavily congested traffic. Blue dots indicate the vehicle moving at considerable higher speeds at the outskirts where the vehicular volume is lower compared to the city center. In the following we discuss the reasons for such a simulation result, and present solutions to them.

3.3.1 Over-comprehensive and bursty traffic demand

The original TAPASCologne O/D matrix yields the traffic demand volume depicted in Fig. 3.4(a), which shows the number of vehicles injected in the whole road network every second. By analyzing the O/D matrix, its source data, and its effect on the microscopic mobility simulation, we identified and fixed the following three problems.

3.3 Repairing the dataset

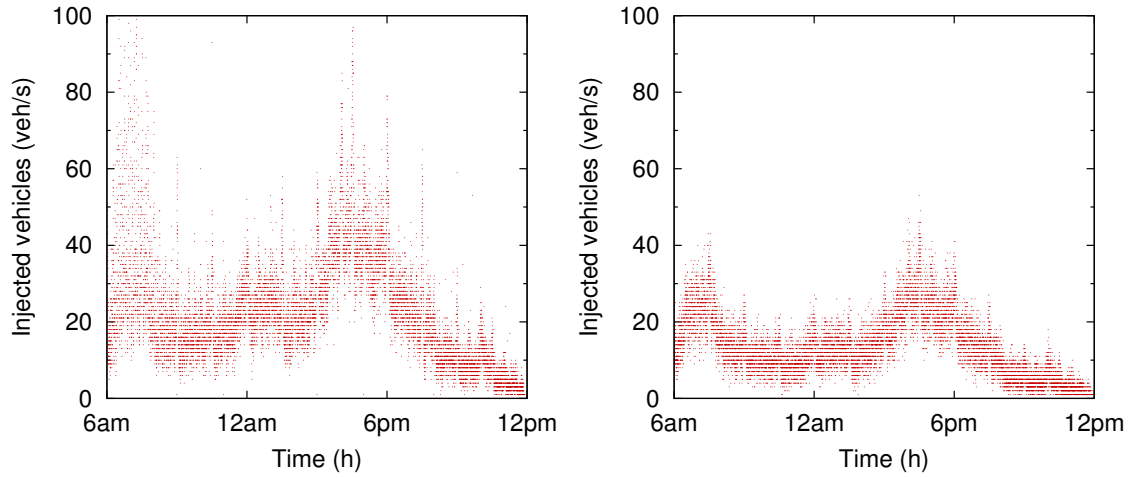

(a) Original

(b) Fixed

Figure 3.4: Volume of traffic injected in the road network between 6:00 am and 12:00 pm, according to the original and fixed TAPASCologne O/D matrix

First, the demand in the O/D matrix is not limited to the vehicular traffic; rather, it includes information on the daily trips of all Koln inhabitants, independently from whether their walk to their destination, or employ public transports, or take a car as either passengers or drivers. Clearly, we are only interested in the latter kind of mobility, since the volume of vehicular traffic directly maps to that of car drivers. According to [HW04, Fig. 4], car drivers account for approximately 50% of the overall trips in the TAPASCologne O/D matrix: thus, we adjusted the O/D matrix by only considering that one trip every two concerns the movement of a vehicle.

Second, the original demand presents an unrealistic variability in the injected traffic over short time scales. This can be observed in Fig. 3.4(a), where, within the span of a few minutes, peaks up to 200 vehicles/s in the injected traffic alternate with instants of reduced injected traffic as low as 10 vehicles/s. Such an excessive burstiness is hardly observable in the reality, especially considering that the injected traffic is aggregated over a very large area. Indeed, we observe these peaks of traffic to be a major cause of congestion, forcing large masses of cars to try entering the road topology at the same time, and thus creating sudden traffic jams. In order to address this issue, we smooth down the original O/D matrix, by adding to the departure time of each vehicle a random offset uniformly distributed in the interval $[-5, 5]$ minutes. This allows to remove the injection bursts, yet retaining the traffic demand properties over larger time scales.

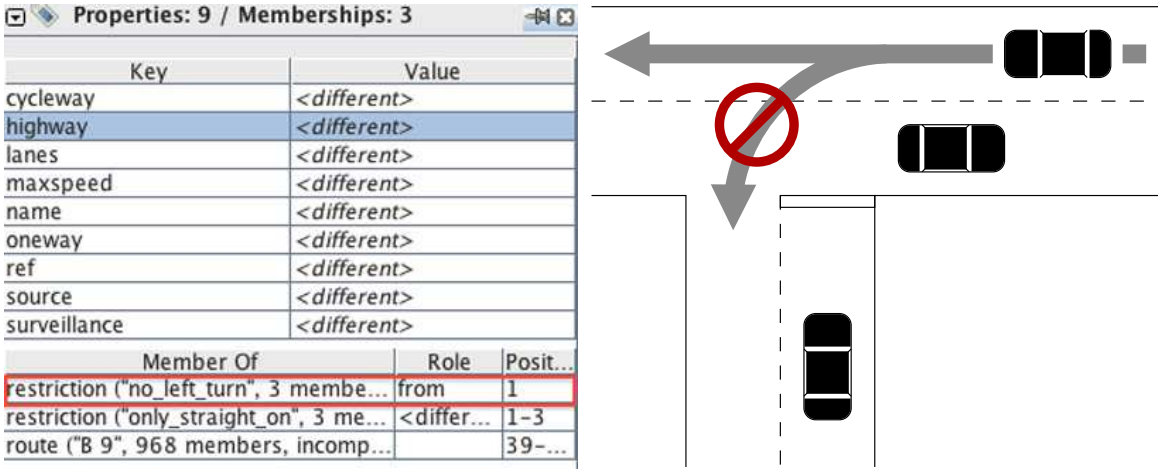

Figure 3.5: Example of wrong restriction enforcement in OpenStreetMap data

Third, the legacy TAPASCologne demand only includes trips starting or ending within the 400 km^2 simulated region. Therefore, vehicles crossing the area entirely, mainly highway traffic around the urban center, are not accounted for. We resort to data made available by the traffic information system of the Nordrhein-Westfalen Ministry of Transport [AUT12] and introduce the missing highway traffic in the repaired demand. The daily evolution of the final traffic demand is depicted in Fig. 3.4(b).

3.3.2 Inconsistent road information

A second source of errors in the simulation was identified in the OSM road data. Although very complete from a topological viewpoint, the OSM map embeds information at times inconsistent with respect to reality. The impact of such inconsistencies, albeit negligible on most of the usages of OSM, can be dramatic for the simulation of vehicular mobility.

A first type of inconsistency is represented by wrong traffic movement restrictions enforced on some road segment. Consider the situation in Fig. 3.5: there, a *no left turn* restriction (bottom of the left image) is present in the OSM road information, for the east-west lane of the horizontal road (right plot). This prevents a car traveling along such a lane to turn left, as in the example in the figure. The OSM data contains at times restrictions of this kind which are not actually there in the real world: the shape of the street layout is not affected by such errors, however the microscopic traffic simulation is, since they can cause vehicles to perform long detours or to get stuck by waiting indefinitely for a possibility to turn and continue their journey. We also identify wrong restrictions by checking the features of congested intersections and T-junctions against the real-world road signalization through the Google Street View service: when necessary, we correct the OSM data according to the

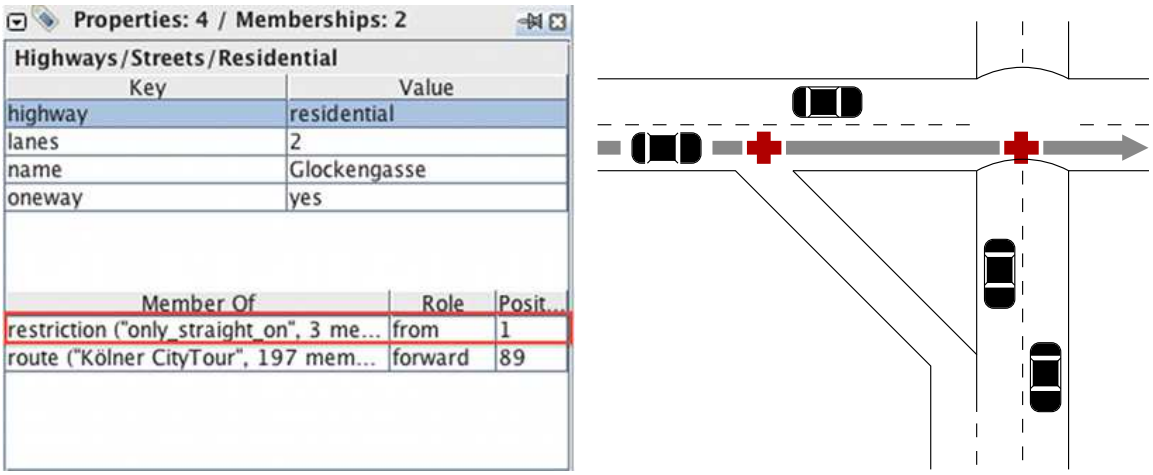

Figure 3.6: Example of continuous restriction enforcement in OpenStreetMap

visual inspection. This allowed us to fix approximately one thousand erroneous restrictions over the road topology in the area under study.

We also identify a second type of inconsistency, represented by correct movement restrictions being enforced on the whole road extent, whereas they only apply to road portions. Fig. 3.6 portrays an example of such a situation, where two one-way roads, going from west to east and from north to south, respectively, cross each other. In the real world, the roads pass one over the other, and vehicles traveling on the horizontal road can join the southbound traffic flow by means of the slanting relief route. In the OSM road representation, the horizontal road is formed by a sequence of segments joined by links (respectively depicted as grey thick lines and red crosses on the right plot); links allow to represent crossings with other roads in the area. The horizontal road is tagged as *only straight on* (left image), a restriction that affects all of its segments: this correctly forces cars to proceed straight at the bridged intersection with the vertical road. However, the same restriction also applies to the previous segments, preventing vehicles from taking the relief route; as a result, the eastbound traffic cannot join the southbound one. Incorrect restrictions of this kind force vehicles to long detours in order to reach their destinations, resulting in a higher traffic volume over the road topology. We identified such situations in most of the interchange nodes among high-speed roads (arterial roads in the city, the freeway ring around downtown Koln, and highways passing close to the urban agglomeration), preventing traffic from correctly switching among such major ways. We solved the problem by separating the segments of a same road and assigning correct restrictions to each of them, repeating the process for approximately 800 roads.

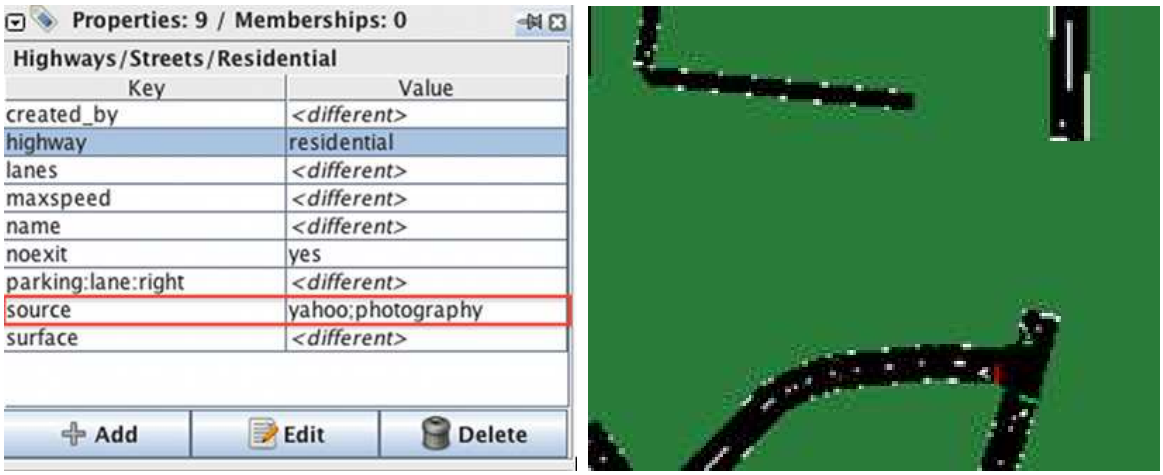

Figure 3.7: Example of unrecognized information ignored during map conversion

3.3.3 Flawed road topology conversion

The OSM road information is natively imported by SUMO through an automated conversion process that proves not to be error-free. A first cause of problems is the presence of road information not recognized by the conversion tool. E.g., attributes with two values are considered as incorrect by the converter, and the associated roads are not included in the topology used for the simulation. An example is shown in Fig. 3.7: the double value of the *source* field (left) causes SUMO not to account for the associated road in simulation (right, a road should connect the north and south branches). We corrected the OSM data so to make all attributes compatible with the SUMO converter.

A second critical aspect is the fact that the topological information in OSM is, at times, simply unfit to be directly converted to the SUMO street layout. An example is depicted in Fig. 3.8. There, the real-world aerial photography of a rather complex intersection (top left), the associated Google map information (top right), and the OSM road topology (bottom left) match. However, the conversion of the latter in SUMO results in an exceedingly intricate intersection, where vehicles get stuck and rapidly form a permanent traffic jam (bottom right). The reason for such a simulation result is that, since two segment links (white dots in the bottom left plot) are present, SUMO interprets the OSM topology as if two road junctions co-existed, one placed right after the other. As a consequence, the number of traffic lights that regulate the car flows into the crossroad is doubled, and yield signs are placed right at the middle of the intersection: the result is the impossibility for vehicular in-flows to correctly merge at the intersection. In order to fix such a problem, we act directly on the OSM road information, by joining road segment links that refer to the same physical intersection. Such an operation allows then a correct conversion by SUMO, so that no traffic jams are observed

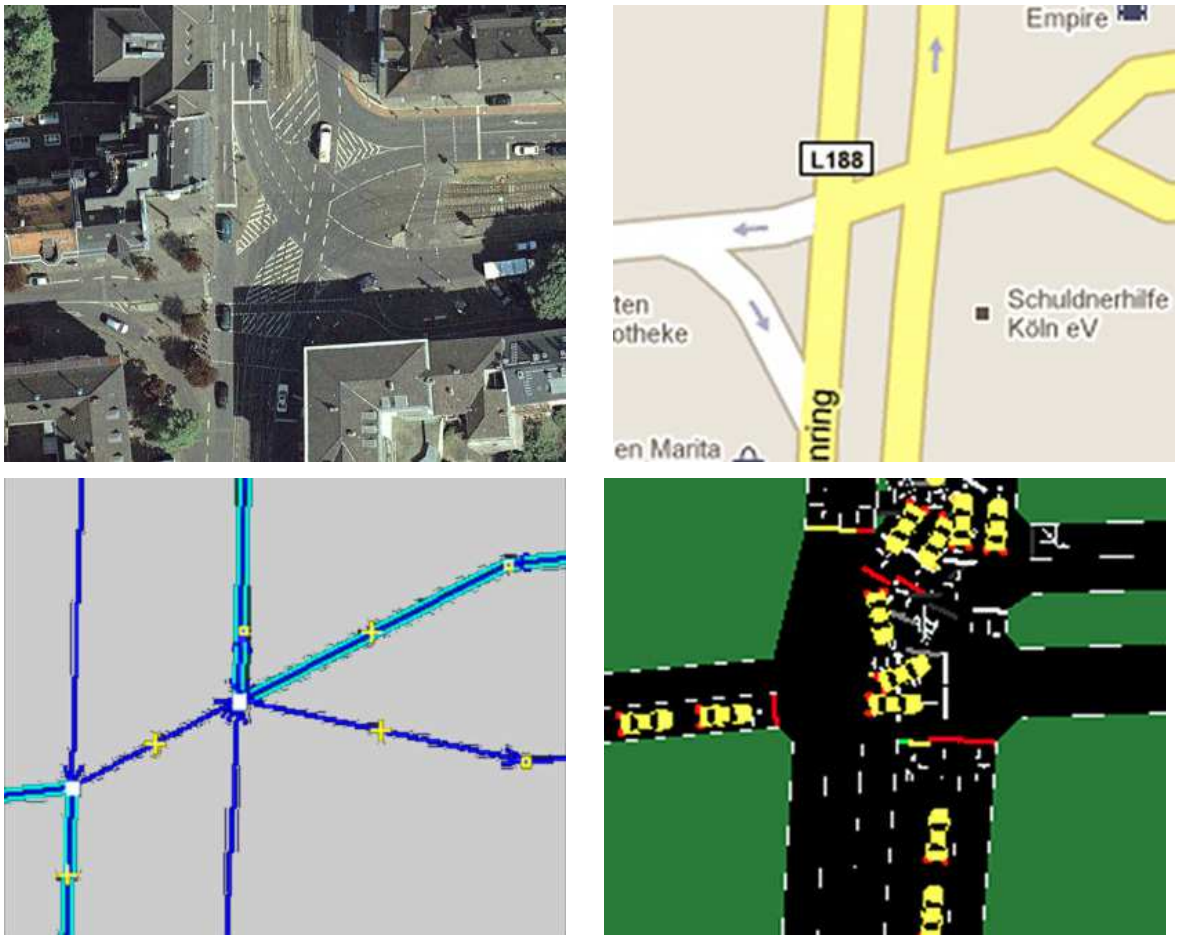

Figure 3.8: Example of topological information unfitness during map conversion

anymore at the road junctions. Additionally, we correct in several cases the number of road lanes entering and leaving an intersection, so as to match the aerial photography data.

The third problem we remarked in the OpenStreetMap-to-SUMO conversion lies in the traffic light deployment. OSM road information already includes data on the presence or absence of traffic lights at road junctions, and SUMO automatically sets the periodicity of the green and red according to the priority of the roads entering each junction. However, the SUMO converter also employs by default a technique to place additional traffic lights over the street layout. After having verified the negative impact of such a traffic light guessing, we disable it. In addition, we identify a number of situations where the presence of traffic lights is not beneficial, and indeed does not correspond to reality: in particular, this is often the case for intersections formed by peripheral roads with identical priority but very unbalanced traffic. Indeed, the similar traffic light periodicity assigned to each road in such a context leads to long queues on the trafficked roads. We thus remove such traffic lights from the OSM data in order to be consistent with the real-world observations.

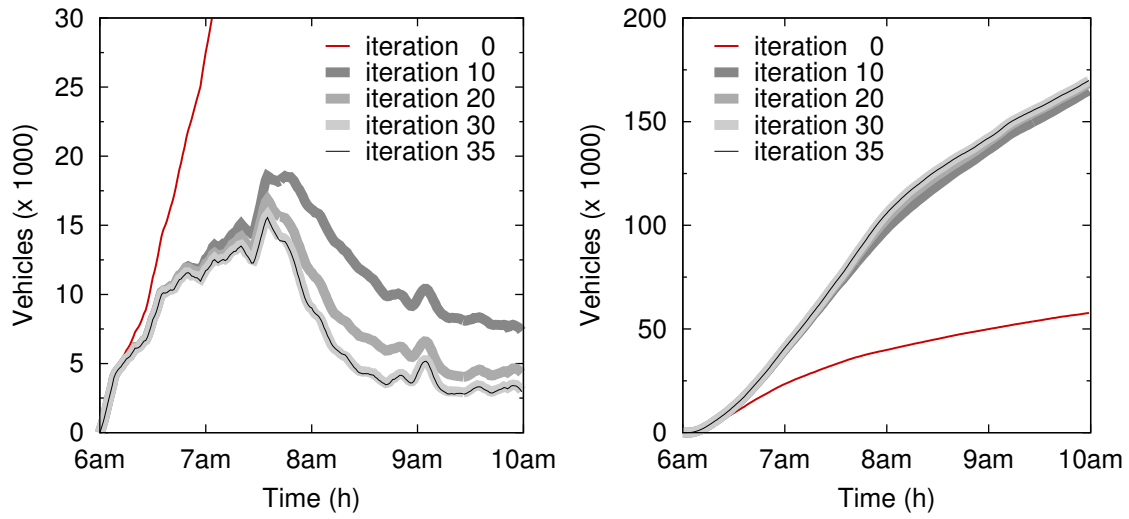

(a) Traveling vehicles

(b) Ended trips

Figure 3.9: Traffic evolution over multiple iterations of the assignment algorithm

3.3.4 Simplistic default traffic assignment

Running the microscopic mobility simulation, with the traffic demand corrected as from Sec. 3.3.1 and the road topology fixed as from Sec. 3.3.2 and Sec. 3.3.3, still results in large congestion and continuous traffic jams all over the street layout. The reason lies in the traffic assignment, i.e., the way drivers choose the route to reach their intended destination. Indeed, SUMO employs a simple Dijkstra's algorithm on the road topology graph, by weighting edges, i.e., road segments, on their length, as well as on the maximum speed they allow: clearly, shorter and faster roads are preferable, and thus are associated with smaller weights. Unfortunately, this means that drivers having similar origin and destination points all choose the same routes for their trips: as a result, they concentrate on major roadways, which are rapidly filled to their maximum capacity, whereas slower or minor roads remain unused. Obviously, high-speed roads alone cannot handle the whole demand in the region, and thus the traffic assignment needs to be improved.

To that end, we resort to the traffic assignment technique proposed by Gawron and presented in Sec. 3.2.4. Such a technique needs to iterate over multiple simulations in order to achieve a dynamic user equilibrium, however the number of iterations cannot be known a priori. Thus, we run the traffic assignment until no significant difference is observed between subsequent iterations.

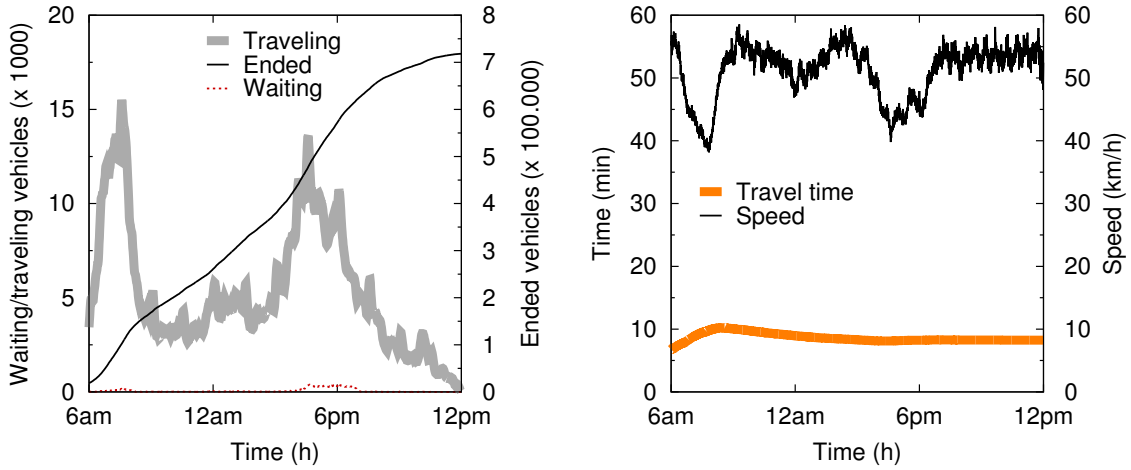

(a) Simulated vehicles status

(b) Mean travel time and speed

Figure 3.10: Final TAPASCologne dataset. Traffic features over time

Fig. 3.9 shows the evolution of traffic while iterating the assignment algorithm. The number of vehicles traveling at the same time over the road topology, in Fig. 3.9(a), tends to explode during the first iterations, as it happens before patching the demand and road topology. However, as Gawron's algorithm iterates, the car traffic is progressively reduced, since drivers tend to employ the different available routes and thus better exploit the capacity of the road network. Fig. 3.9(b) confirms that iterations significantly improve the traffic conditions, as they increase the number of vehicles that reach their destination, successfully ending their trip. Similar trends are observed for the other traffic metrics, and in all cases iterations after the 35th do not produce any noticeable improvement. Therefore, in the following, we consider the traffic assignment obtained at the 35th iteration.

3.4 Koln vehicular mobility dataset

The resulting dataset comprises nearly seven hundred thousand trips of vehicles in the Koln larger metropolitan area, over a period of 24 hours which makes our dataset comparable with the Zurich region trace. The simulated traffic now mimics the normal daily road activity in the region, as the fixed road topology can accommodate the updated traffic demand and assignment. Fig. 3.10(a) shows that the vehicular volume in our simulated dataset follows a typical day real-world measurement in any city (e.g., [Cor13]) which peaks to its highest during the traffic rush hour. By comparing it to the equivalent plot before repair, in Fig. 3.2(a), it is clear that the number of traveling cars now follows the traffic demand, with peaks during

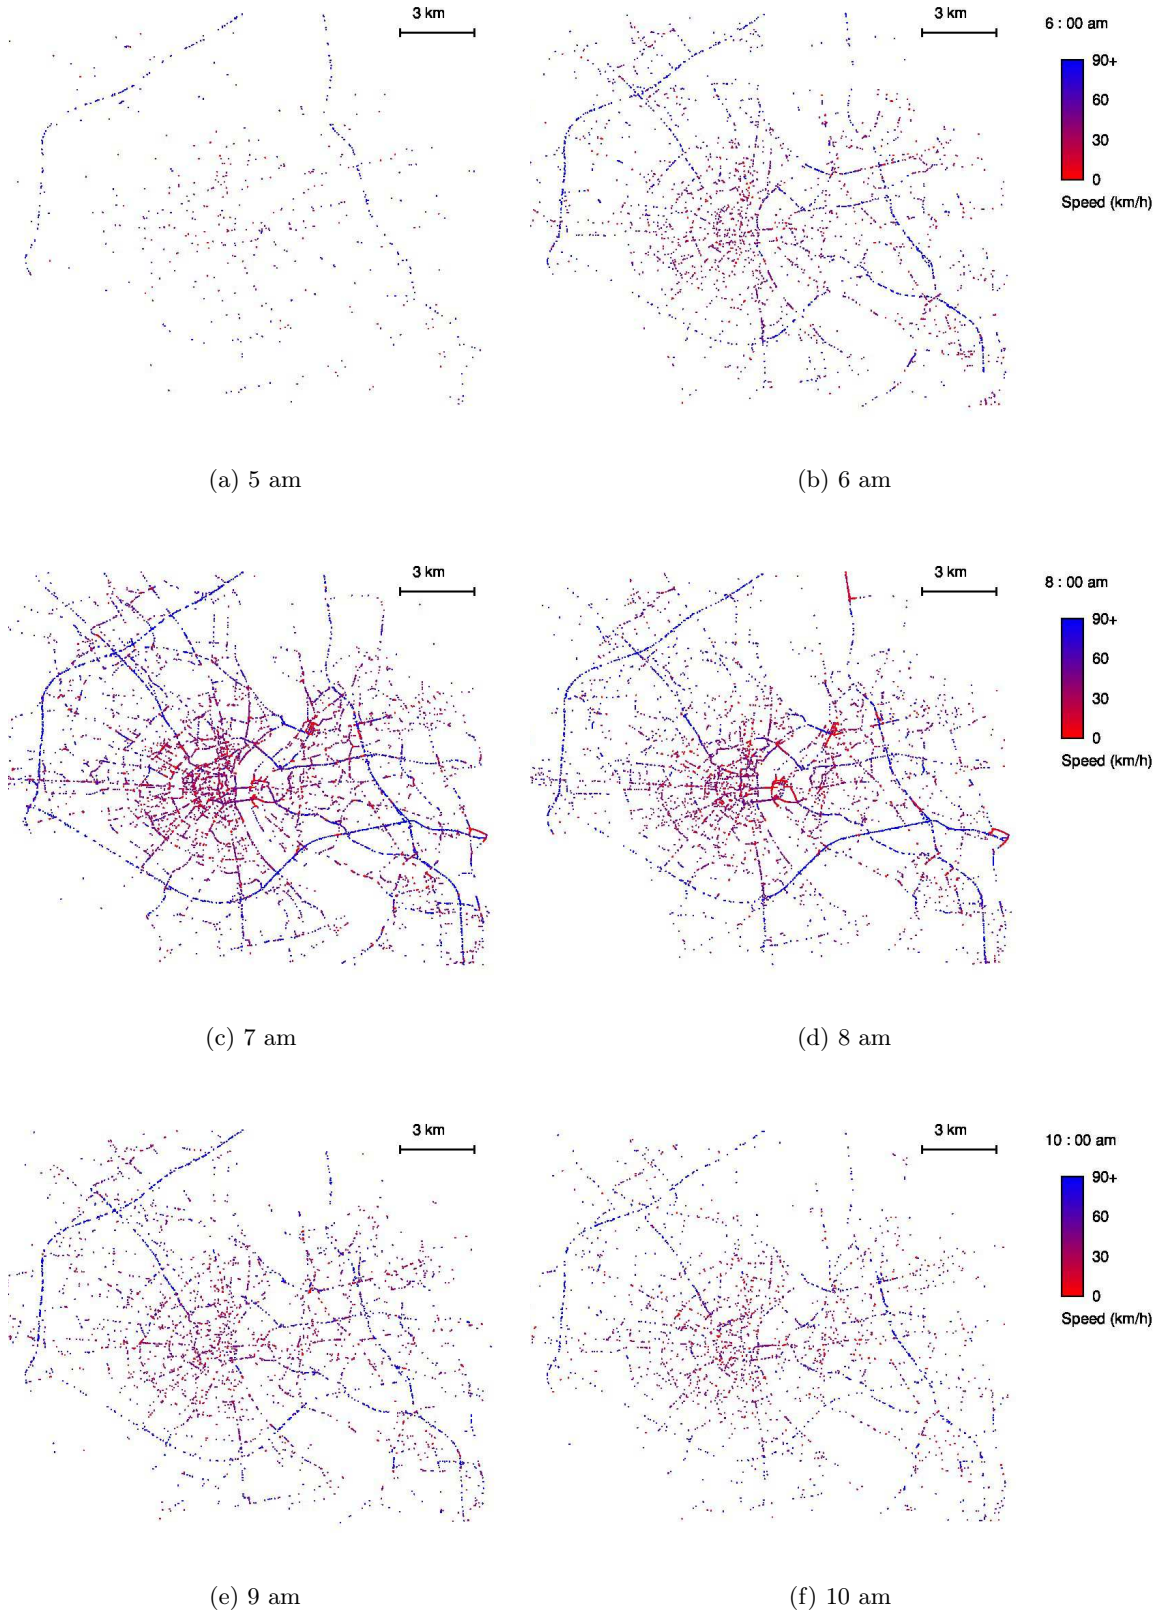

(a) 5 am

(b) 6 am

(c) 7 am

(d) 8 am

(e) 9 am

(f) 10 am

Figure 3.11: Spatiotemporal distribution of vehicular traffic density (5 am to 10 am).

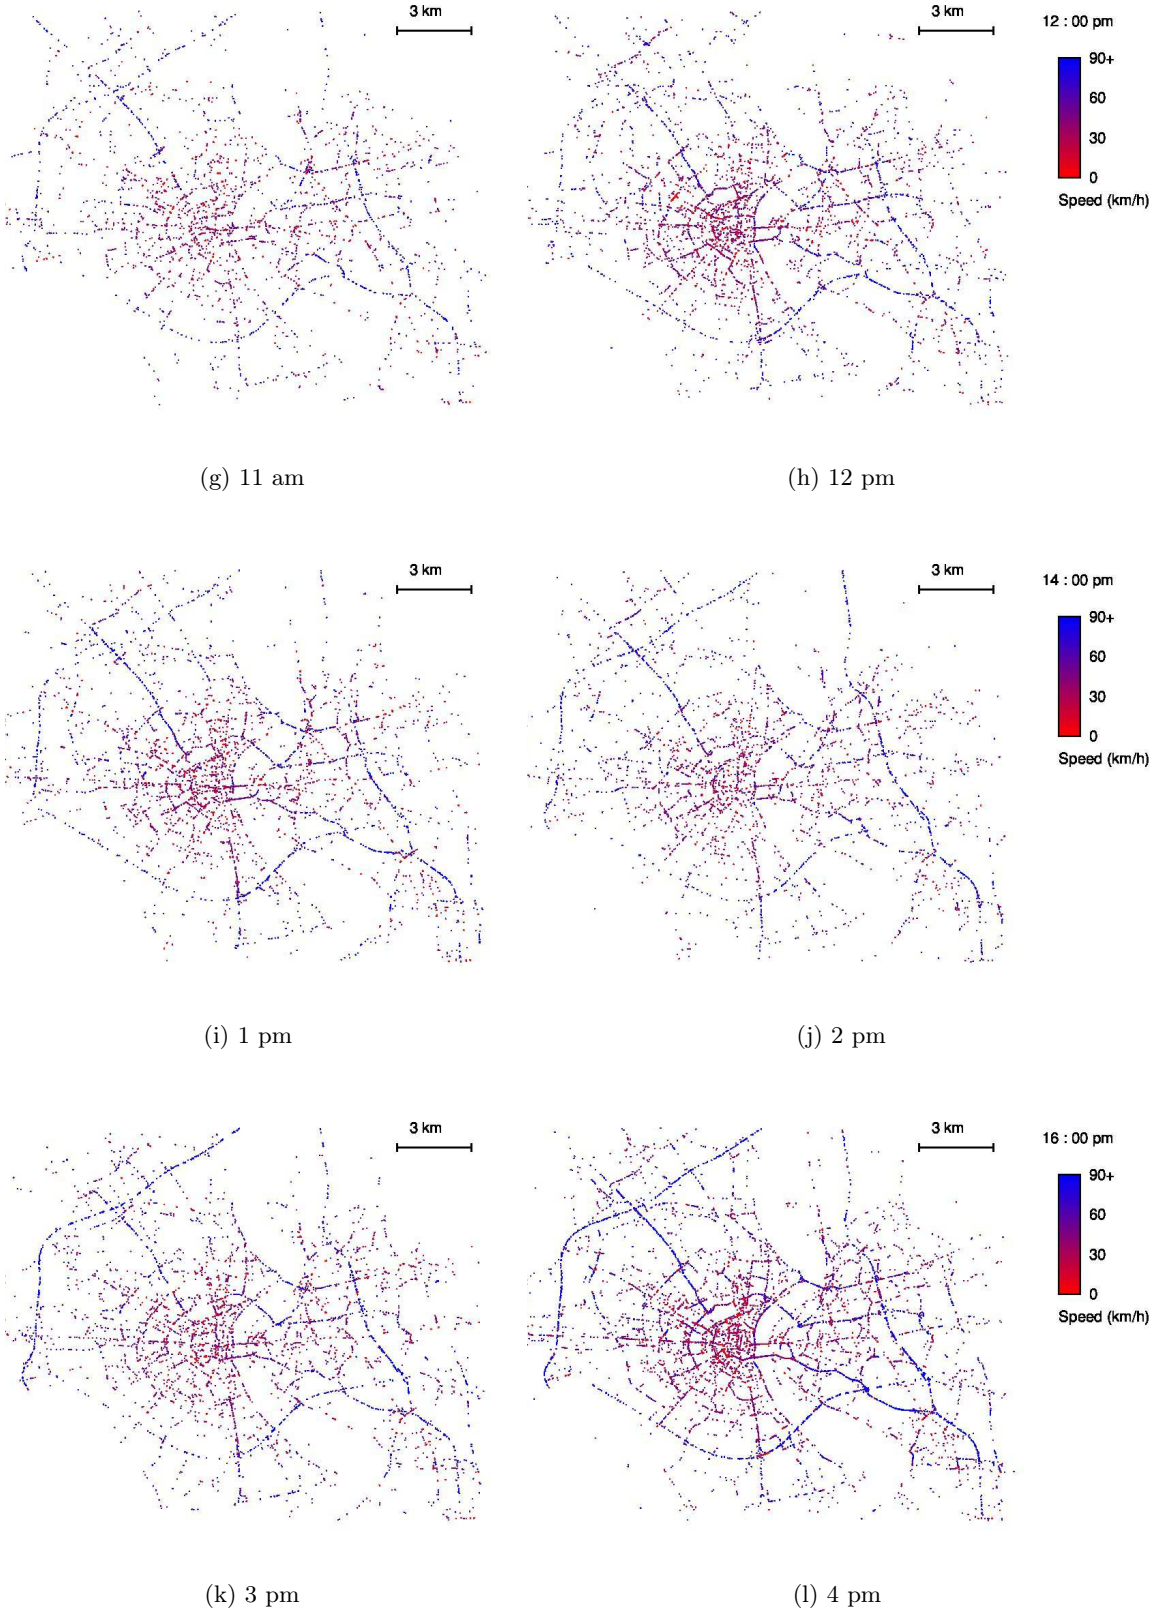

(g) 11 am (h) 12 pm

(i) 1 pm (j) 2 pm

(k) 3 pm (l) 4 pm

Figure 3.11: Spatiotemporal distribution of vehicular traffic density (11 am to 4 pm).

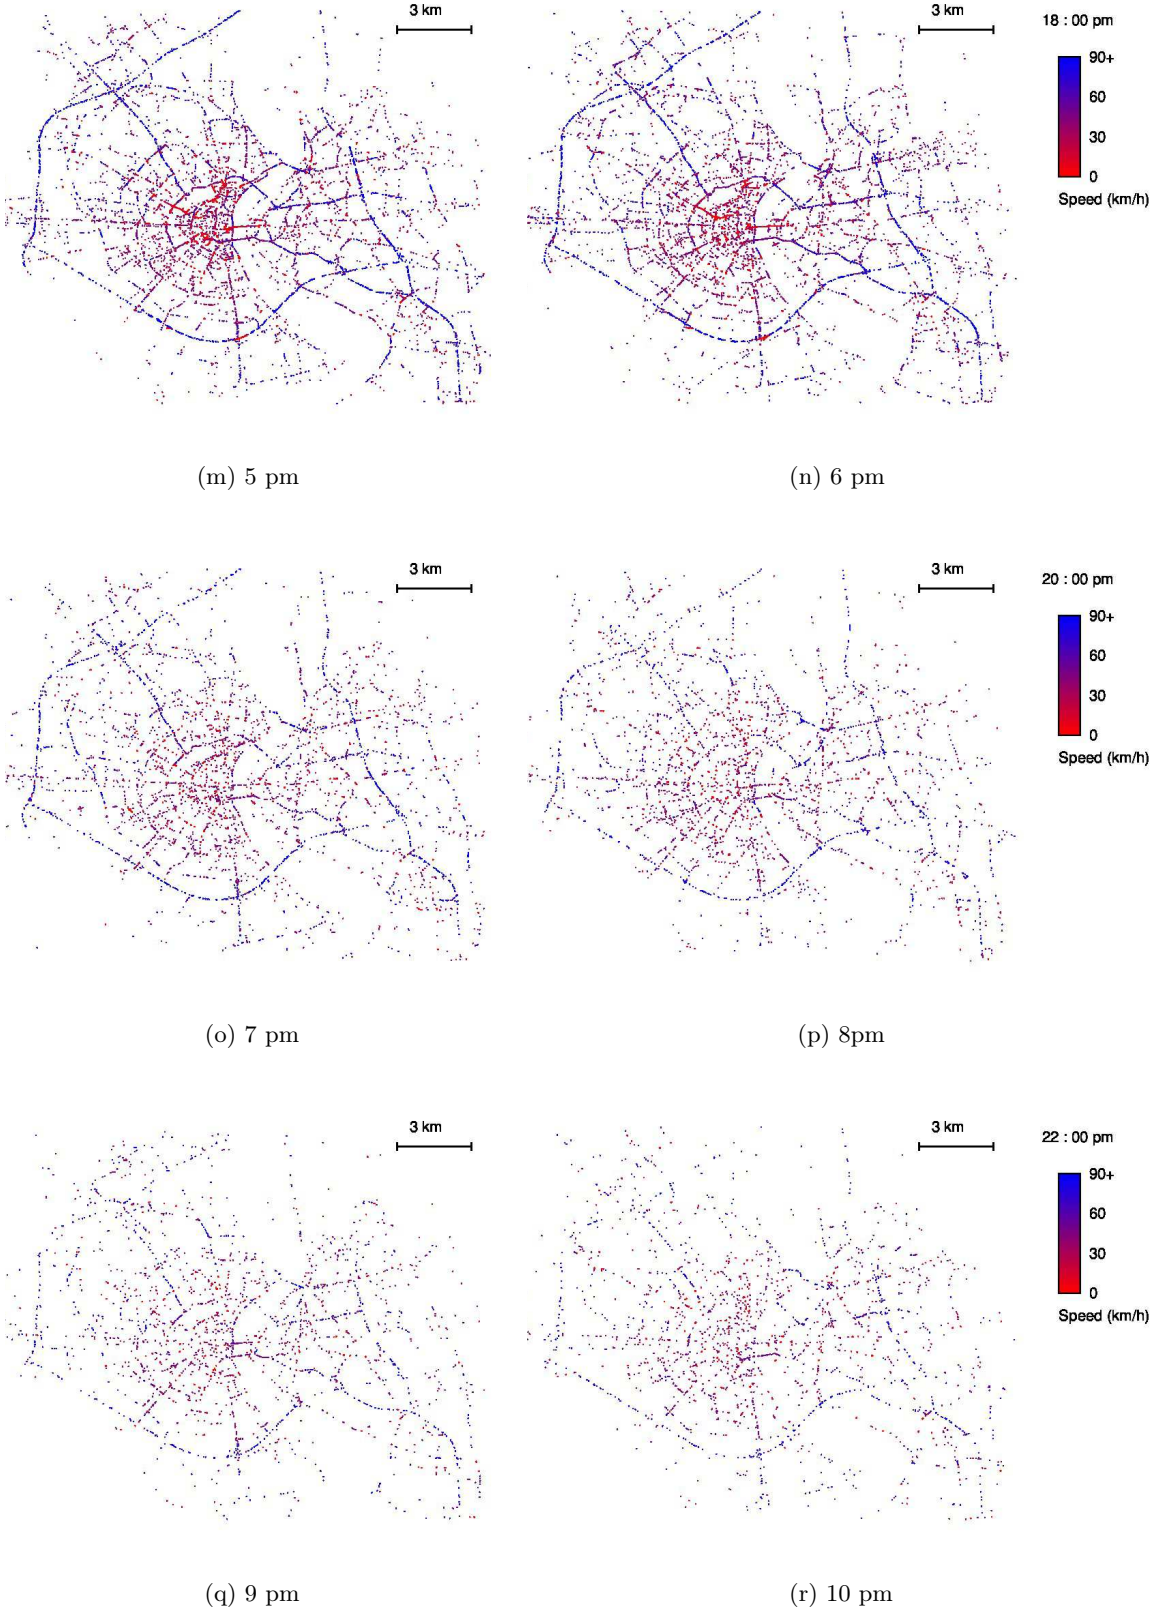

(m) 5 pm

(n) 6 pm

(o) 7 pm

(p) 8pm

(q) 9 pm

(r) 10 pm

Figure 3.11: Spatiotemporal distribution of vehicular traffic density (5 pm to 10 pm).

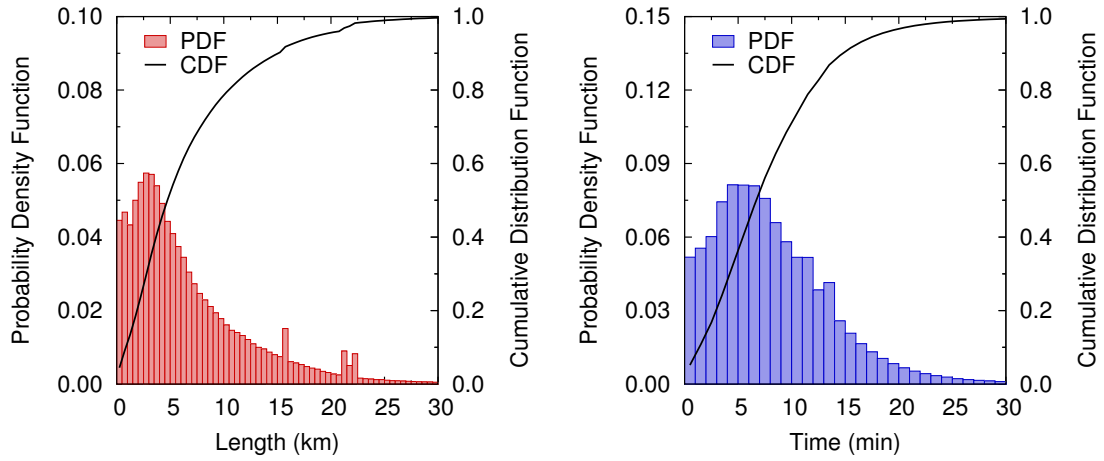

(a) Length

(b) Duration

Figure 3.12: Final TAPASCologne dataset. Trip feature distributions

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 15000 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. This features out-stands our dataset with that of Zurich region, where the off-peak vehicular mobility is missing despite realistic peak hour traffic behaviors.

Also, the number of ended trips now 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. The average travel time and speed recorded during the morning, in Fig. 3.10(b) confirm the previous results, as we observe quite constant behaviors, only modified during the peak hours.

As a result, the road traffic evolution over time is achieved and is presented in Fig. 3.11. Following Fig. 3.10(a), It is obvious that vehicular density at off-peak hours like 5 am - 6 am, 10 am - 2 pm and 8 pm onwards are relatively low when compared to the peak time at 6 am - 8 am and 3 pm - 6 pm. Large portions of the urban roads are in violet, indicating fluid traffic conditions. Fewer dark red regions exist during the off-peak hours because of the fact that vehicles follow the traffic rules and are obliged to lower their speed or stop at traffic signals. But these dark red regions grow in peak hours in the city center, time when most people travel to work in different part of the urban region but at some point in time follow the same route thereby leading to traffic congestions with slower speeds below 50 km/h. Similar behavior is seen during the evening peak hours when people travel back from office to home or

to leisure activities. On the other hand, vehicular speeds at highways are denoted by a bright blue, corresponding to speeds higher than 90 km/h. This is because of flexible speed limit and fewer or no traffic lights condition along the highways. Such detailed traffic density dynamics cannot be obtained in Zurich region dataset as it is limited only to highways and freeways. Other datasets which deliver great detailed are limited by duration of the dataset or extends over a smaller geographic area like Turin and Karlsruhe. Comparing the road traffic at 7:00 am, in Fig. 5.25(c), looks significantly better than the original one, in Fig. 3.3. These results indicate that the repaired road network with the macroscopic traffic definition discussed in Sec. 3.3 accommodate the traffic load of a typical day leading to a realistic behavior that one could expect in the real-world.

In Fig. 3.12, we analyze some interesting features of the trips in the dataset. The length of routes traveled by drivers in the simulated region, in Fig. 3.12(a), appears to be mostly in the order of a few kilometers, as one would expect in a urban area. Yet, trips longer than 10 km are not uncommon, as one driver over five has to travel such a long road: those are mostly commuters living in the suburbs of Koln. Also, note the peaks around 15 and 20 km: these are the contributions of the traffic crossing the whole region, whose route length is thus constrained to the exact length of the highway segment in the simulated area. The route length has a clear impact on the distribution of trip durations, in Fig. 3.12(b), where we find the same shape of Fig. 3.12(a) and we can observe how half of the trips last less than 8 minutes.

Interestingly, we found the macroscopic traffic simulated in the final TAPASCologne dataset to nicely match that observed in the real world, through real-time traffic information services. In Fig. 3.13, we compare the simulation output, at 5:00 pm, with the road traffic information retrieved through the ViaMichelin and GoogleMaps live traffic services at the same hour. This represent 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. Although we acknowledge that more rigorous tests are needed to fully prove the realism of the dataset, we regard the result as very encouraging, especially considering that more complex assessments are unfeasible at this moment due to the unavailability of sensible data (such as that provided by, e.g., traffic counter records). Comparative analysis with real-world live traffic is the first in its kind where other datasets fail to provide either because of low detail macroscopic model adopted which does not mimic the real-word live traffic or due to lower number of traffic which they consider.

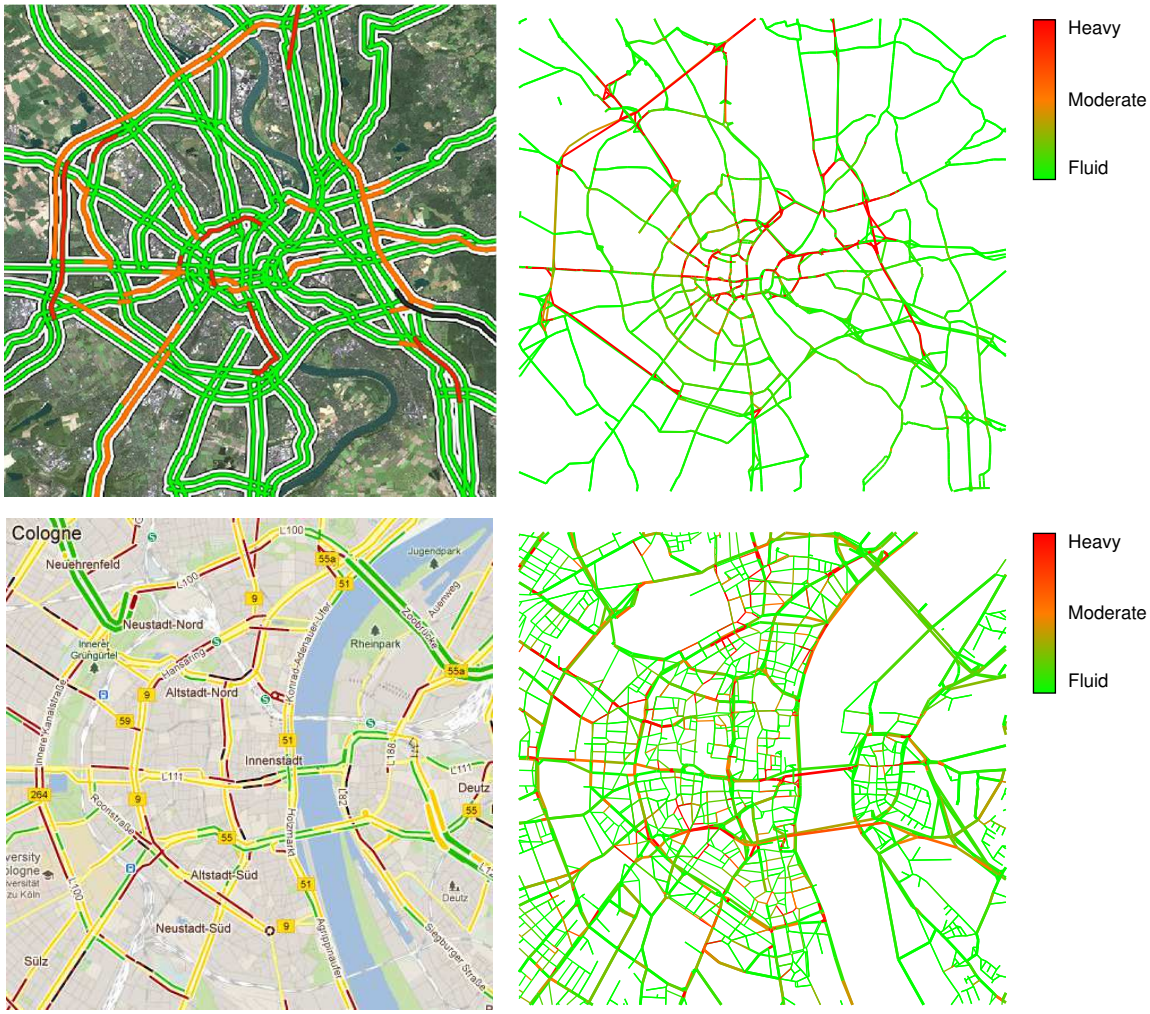

Figure 3.13: Final TAPASCologne dataset. Traffic at 5:00 pm, in the real world (left) and in simulation (right). The plots refer to the whole urban area (top) and in the city center (bottom). This figure is best viewed in colors.

Referring back to Sec. 2.5.6 and comparing the features of well know datasets discussed in Sec. 2.5, the only dataset that seems to meet all the requirements is the Koln dataset. It couples a large scale with a high level microscopic detail and an accurate macroscopic traffic description.

3.5 Integrating the Koln dataset into OpenAirInterface

Koln dataset has been made available in formats which are widely accepted in various network simulators. This gained it to penetrate well in to the telecommunication research community. One of the unique work involves porting the mobility dataset to a LTE emulator known as *OpenAirInterface* (OAI) [EUR12].

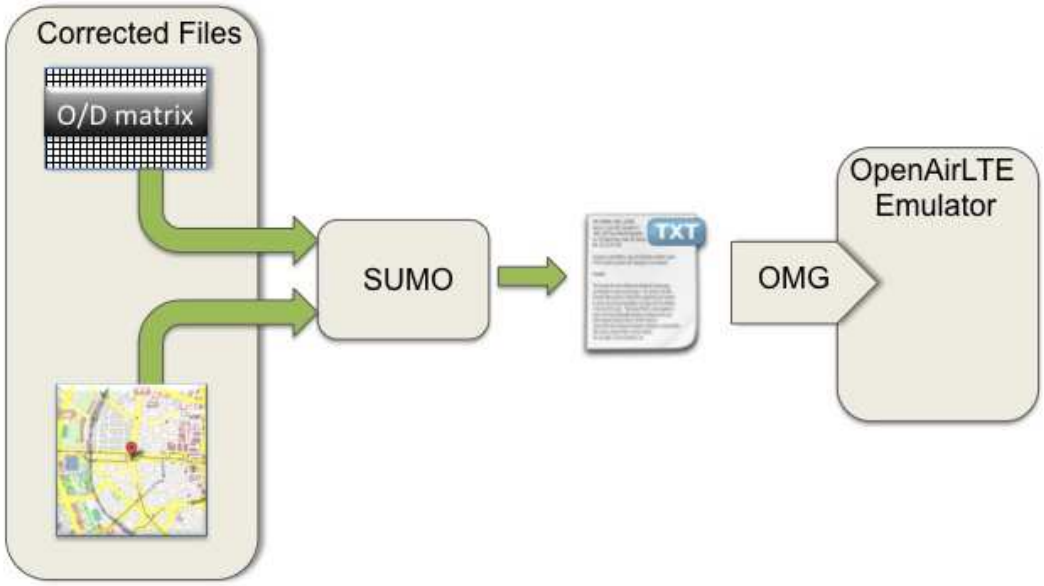

Figure 3.14: Static mobility feed from SUMO to OpenAirLTE.

OAI is an open-source hardware/software development platform and open-forum for innovation in the area of digital radio communications. It was created by the Mobile Communications Department at EURECOM. The emulator platform permits to configure environment, network topology, application traffic, and emulation IO parameters. The behavior of the wireless medium is modeled using a physical layer (PHY) abstraction which emulates the error events in the channel decoder and provides emulated measurements from the PHY in real-time. The communication between the system entities i.e., eNB's and UE's involves initialization and synchronization in real time. In such a setting, mobility of these entities in the emulation environment follows a subset of the Koln mobility dataset. This was made possible with two models, namely *Static* and *Real-time*.

Fig. 3.14 shows the static model where the SUMO simulation logs the vehicular node mobility in a text format, OpenAirInterface Mobility Generator (OMG) module reads this mobility logs to define the UE's mobility (assuming vehicular node as UE). On the other hand, Real-time model as depicted in Fig. 3.15 involves real time mobility information exchange between the SUMO simulator and the OAI emulator's OMG module. This communication is made possible by a re-implemented version of Traffic Control Interface (TraCI) [WPR+08], a technique for interlinking SUMO with external simulators. Along the simulation, the updated vehicular node position is conveyed to OMG periodically, which intern updates the node position in the emulation. Such studies involving realistic platforms in large scale are of great significance for protocol validation, testing and performance analysis in future wireless communication.

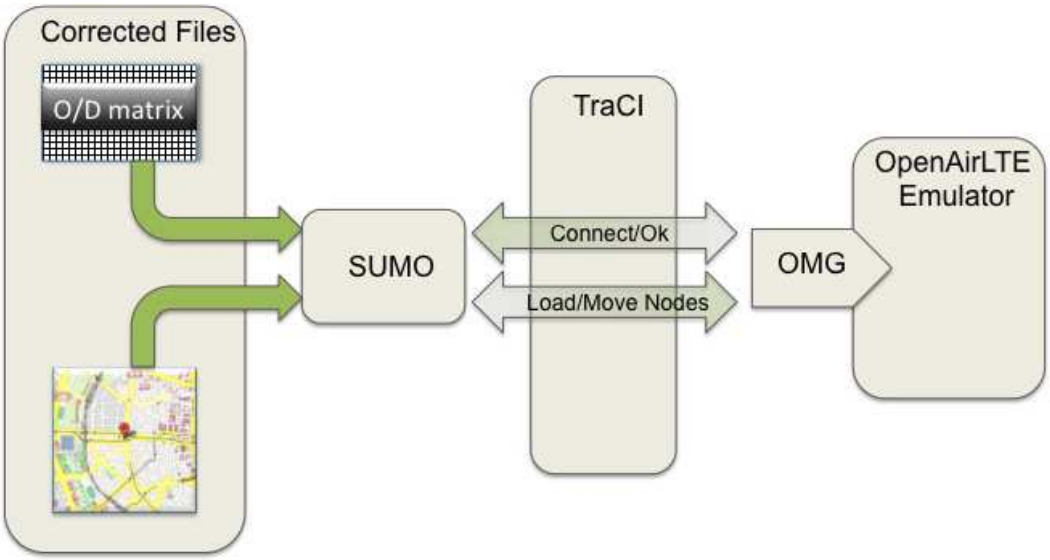

Figure 3.15: Real time mobility feed from SUMO to OpenAirLTE through TraCI.

3.6 Summary

In this Chapter, we proposed a novel, highly realistic vehicular mobility dataset. We discussed the various tools that are employed in its generation. The issues and the solutions adopted to address integration of these tools were presented. The resulting dataset was found to follow a typical day vehicular flow in any metropolitan city over a 24 hour time period. In addition, when mapping the dataset with the live traffic provided by ViaMichelin and GoogleMaps, we found traffic congestions in the same areas of the city as in the real world. To the best of our knowledge such realism was not provided by any other dataset to date. Later, integration the dataset to an LTE emulator is discussed. The dataset proposed in this Chapter has been widely adopted by the research community worldwide: the dataset has been downloaded more than 700 times from more than 40 different countries since November 2011. Research teams those using our dataset include: KIT, Germany; EURECOM, France; IMDEA Networks, Spain; Universidade Federal de Minas Gerais, Brazil; Ilmenau University of Technology, Germany; University of Luxembourg; Tampere University of Technology, Finland; University of Cyprus, Cyprus.

Finally, this work also exposed the flaws in the sumo NETCONVERT module that is responsible from the conversion of the real world map from OpenStreeMap service into the simulation. This led to the enhancement of an option known as *junction.join* in SUMO simulator.

Impact of user mobility in wireless networks

4

Contents

4.1 Introduction 58
4.2 Access network 58
4.2.1 RAN infrastructure deployment 58
4.2.2 Macroscopic analysis 59
4.2.3 Microscopic analysis 68
4.3 Autonomous network 77
4.3.1 Modified Koln dataset 77
4.3.2 Vehicular contact 78
4.3.3 Network graph metrics 79
4.3.4 Epidemic dissemination 84
4.4 Summary 87

4.1 Introduction

Hindered by the detailed realistic vehicular mobility information in large metropolitan areas, little information is available on the dynamics of real-world like road traffic and on their relevance to networking solutions. The vehicular mobility dataset we introduced in the previous Chapter allows the study of solutions for network planning, network management or vehicular data sharing for future wireless communications in large-scale urban area with a level of detail that was not possible before.

In this Chapter, first we propose a novel study on macroscopic and microscopic features of pervasive vehicular access in a case-study large-scale urban environment with realistic vehicular dataset of the road traffic and Radio Access Network (RAN) deployment. As discussed in Chapter 1 growing penetration of smart devices as in-vehicle component signifies that the vehicles on road, in near future, will be the source and destination of large amount of data traffic inviting additional challenges to the cellular operators. Hence this study gains significance as it unveils the potential areas in communication architecture which need to be revisited to accommodate the anticipated growth of pervasive vehicular access. Second, we quantify the effect of vehicular dynamics with a V2V perspective to understand the level of network robustness that can be derived and also analyze how a network application reacts to the level of realism brought about by our Koln dataset.

4.2 Access network

In this section, we introduce the RAN infrastructure geographic deployment in Koln metropolitan region and the procedure we adopt to mimic the cellular network like coverage. With such a setup we analyze the macroscopic and microscopic features exhibited by the vehicular mobility on the road network in time and space.

4.2.1 RAN infrastructure deployment

Information on the real-world deployment of the cellular access infrastructure in the Koln region was retrieved from a dedicated database [FG12]. We focused on the 247 base stations of one mobile telecommunication operator, whose locations are shown in the left plot of Fig. 4.1. As one can expect, the base station deployment is denser at the city center and becomes sparser as we move towards the outskirts, accordingly to the diverse capacity requirements of such areas. In order to approximate the coverage of individual base stations in the region, we employed a Voronoi tessellation based on the base station locations. Such an approach

Figure 4.1: Deployment of the RAN infrastructure in Koln, for one mobile operator. Base station locations (left), and cell layout resulting from the Voronoi tessellation (right).

is commonly adopted in the context of mobile data traffic analysis [PSBD11]. The resulting cell layout is portrayed in the right plot of Fig. 4.1. In the present day cellular network, apart from the regular deployment an additional hierarchy of cellular cells are used to cover multiple lower hierarchy small cells and are termed as *umbrella cells*. These umbrella cells are used to ease the frequent call handoff among the smaller cells resulting from the high-speed user mobility or to provide service when a small cell fails. Since we lack the deployment information of such umbrella cells we limit our discussion to the hierarchy in figure Fig. 4.1.

4.2.2 Macroscopic analysis

We first characterize the pervasive vehicular access in the Koln region from a macroscopic viewpoint, i.e., aggregating information about the road traffic over the whole 400 km^2 geographical area. We first study the spatiotemporal distribution of the load that vehicular users are expected to generate on the RAN. This provides us with a global overview of the impact that pervasive vehicular access will have on the RAN deployment. Then, we observe the dynamics of large-scale flows of vehicular users through the metropolitan region, which lets us comment on the inter-cell mobility they induce.

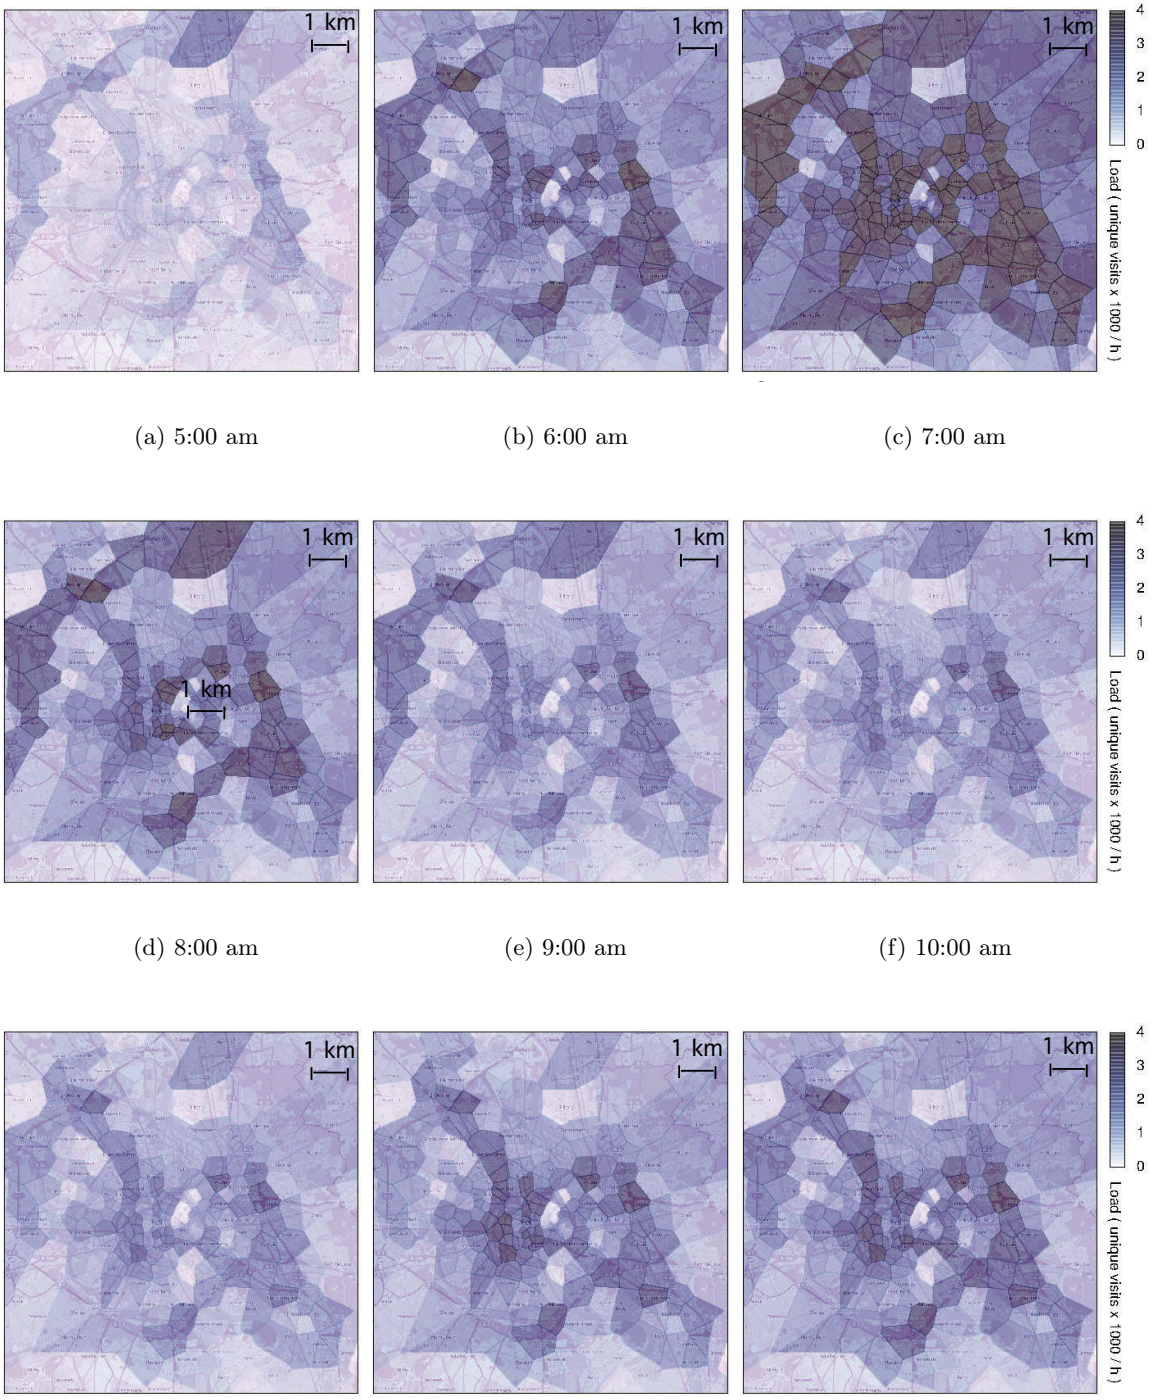

(a) 5:00 am (b) 6:00 am (c) 7:00 am

(d) 8:00 am (e) 9:00 am (f) 10:00 am

(g) 11:00 am (h) 12:00 am (i) 1:00 pm

Figure 4.2: (a)–(i): spatiotemporal evolution of the expected data traffic load generated by pervasive vehicular access in the Koln region, during a typical day (5 am to 1 pm).

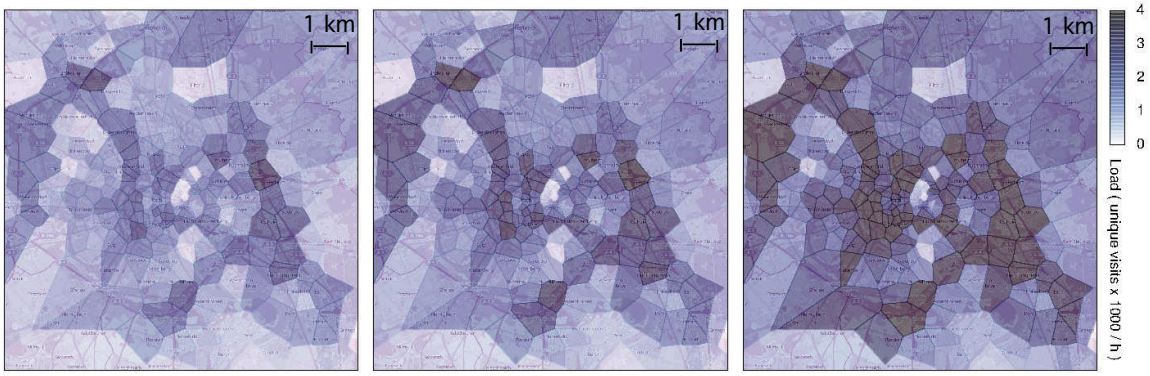

(j) 2:00 pm (k) 3:00 pm (l) 4:00 pm

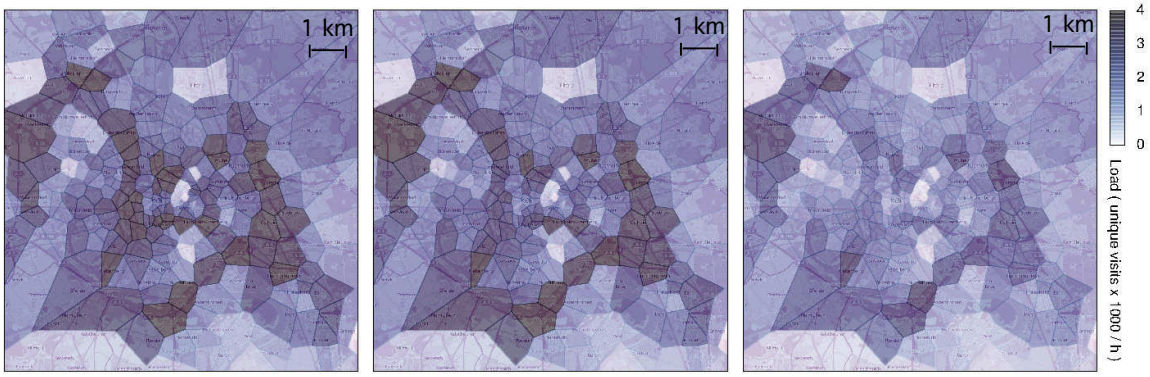

(m) 5:00 pm (n) 6:00 pm (o) 7:00 pm

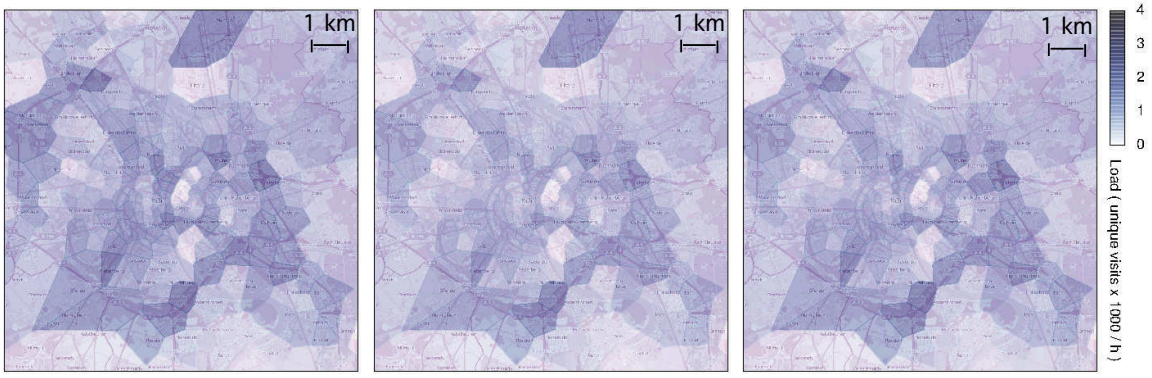

(p) 8:00 pm (q) 9:00 pm (r) 10:00 pm

Figure 4.2: (j)–(r): spatiotemporal evolution of the expected data traffic load generated by pervasive vehicular access in the Koln region, during a typical day (2 pm to 10 pm).

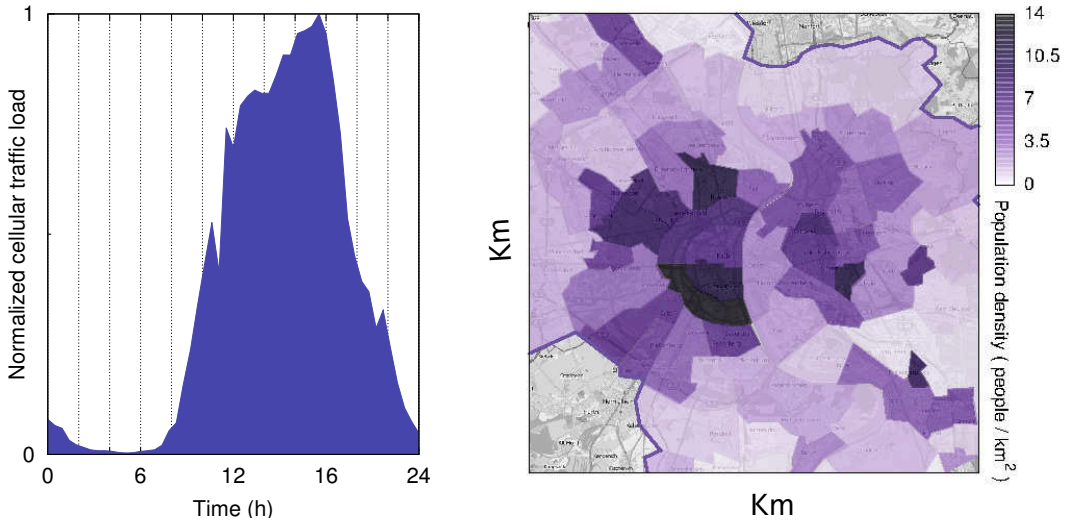

(a) RAN utilization

(b) Population

Figure 4.3: (a): typical daily profile of normalized RAN traffic load – concession of the Autonomous Networks Research Group at USC, `http://anrg.usc.edu`. (b): geographical distribution of the population in the Koln region.

RAN access load

We assume that the communication load expected to prompt on the RAN, is the exact quantity of vehicles passing within the Voronoi coverage region of each base station every hour. Rather intuitively, this choice is motivated by the fact that a higher number of traversing cars yields a greater burden on the base station, in terms of both control signaling overhead and mobile data traffic.

The spatiotemporal evolution of the pervasive vehicular access load is portrayed in Fig. 4.2. Plots refer to different hours of the day, between 5:00 am and 11:00 pm. In each of image, Voronoi cells are colored according to their associated traversing volume, expressed in vehicles per hour, according to the color range at the right of each row. Therefore, light colors denote low vehicular access, while darker colors indicate an increasingly higher load induced by vehicles on the RAN infrastructure.

We can 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 urban area (see Fig. 3.13). 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, the heaviest pervasive vehicular access load perfectly maps onto the highways that surround the urban region as well as onto 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. A somewhat higher access load is observed around midday, but the second daily peak is recorder 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. 4.2(c)) the heavy traffic conditions force drivers to optimize their travel time by spreading more evenly over the street layout. After 7:00 pm, the road traffic activity then progressively decreases, so does the vehicular access.

We can remark the following features of interest from a RAN access viewpoint. First, the pervasive vehicular access shows a *strong daily variability* in space, 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. 4.3(a). 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. 4.3(b), we can observe that the geographical distributions of customary mobile traffic load and pervasive vehicular access do not overlap at all. This implies that the conventional *resource allocation in the RAN* used to serve mobile phone users need enhancements if large number of vehicular users are to be served. 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. 4.2 outline how the relative behavior of *low- and high-load areas remain the same* over the whole 24-hour period.

Inter-cell mobility dynamics

The previous section provided a static view of the spatial distribution of pervasive vehicular access at different time instants. Considering the flows of road traffic grants a different perspective, unveiling the correlation of vehicular use mobility over the Koln region. Specifically, we are again interested in macroscopic vehicular flows, i.e., large groups of cars traveling along similar paths. In order to identify such flows, we analyze the Koln dataset and compute, for each RAN cell in the region, how vehicles leaving it distribute among other cells. In other

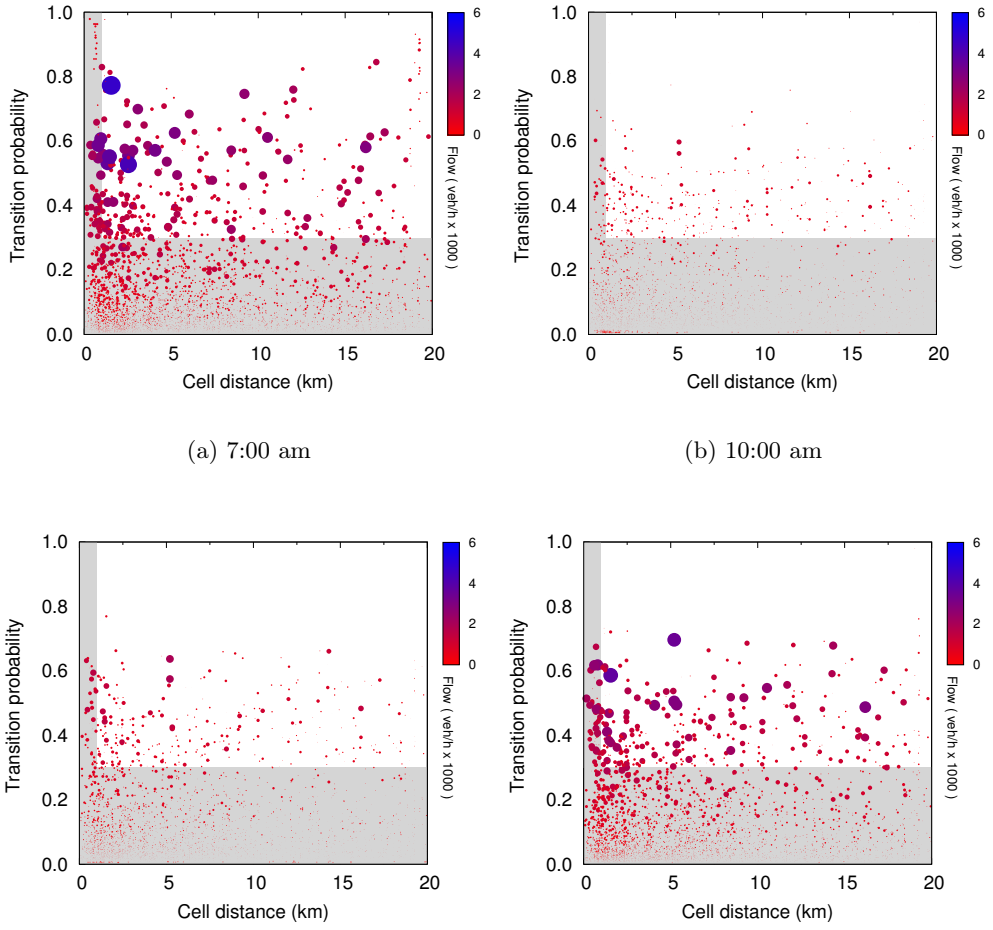

(a) 7:00 am

(b) 10:00 am

(c) 1:00 pm

(d) 4:00 pm

Figure 4.4: Macroscopic vehicular flows in the Koln dataset. Plots refer to four different hours of the day. This figure need to be viewed in colors.

words, we compute the volume of vehicular users that, leaving a first cell, reaches every other cell in the region.

An overview of the resulting macroscopic vehicular flows is provided in Fig. 4.4. The four plots refer to representative times of the day, namely 7:00 am (morning traffic peak), 10:00 am (morning off-peak), 1:00 pm (midday minor peak), and 4:00 pm (afternoon rush hours). In each plot, a single dot represents one inter-cell flow, i.e., the volume of vehicular users leaving a cell i and reaching a cell j. The abscissa of each dot represents the geographical distance between cell i and cell j. The y value is the transition probability from i to j, i.e., the probability that a vehicle leaving i ends up in j during its trajectory. The size and color of the dots are used to describe the total volume associated to the flow, i.e., the number of vehicles per hour that move from i to j. We associate larger flow of vehicles between two cells

that are separated by lower geographic distance by blue dot with greater diameter located close to the origin of the plot.

In all plots flows are mostly characterized by short distances, as dots are denser towards low x-axis values. Such an effect appears especially for large-volume flows. This is expected, since vehicles leaving one cell necessarily move to an adjacent cell, generating a large number of high-volume flows between close-by cells. More interestingly, there exist macroscopic flows of thousands of vehicular users per hour among cells whose geographical distance is in the order of kilometers. Most such flows are also characterized by a high transition probability, implying that a significant portion of vehicular users in the origin RAN cell share a similar mobility pattern.

The observations above yield different levels of significance depending on the day time. As a matter of fact, by comparing the four plots it appears evident that flows between distant cells and characterized by large volumes and high transition probability are much more present during the morning and afternoon traffic peaks. In addition, the dot pattern does not remain the same at different times of the day, although some slight correlations seem to exist. The fact that the position, size and color of dots tend to change among the plots implies that the macroscopic vehicular access flows are not stable over the 24 hours.

In order to better investigate this phenomenon, we plot the geographical localization of the large vehicular access flows identified in the plots of Fig. 4.4. More precisely, we focus on flows that occurs 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). The flows we will consider in the following lie in the white area in the plots of Fig. 4.4, whereas the flows we discard fall in the grey area.

The results of the geographical representation of the vehicular access flows are displayed in Fig. 4.5, 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 Koln urban area, green arrows represent flows moving away from the center, and blue arrow 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.

Early in the morning, i.e., at 5:00 am or earlier, there is small traffic in the region, and thin flows are detected over the highways that surround the urban area (blue arrows in Fig. 4.5(e)) or along the primary roads that bring inbound traffic towards the city center (red arrows in Fig. 4.5(e)). The mobility pattern does not change later on, but the volumes dramatically increase. Between 6:00 and 9:00 am both peripheral highway traffic and the inbound traffic

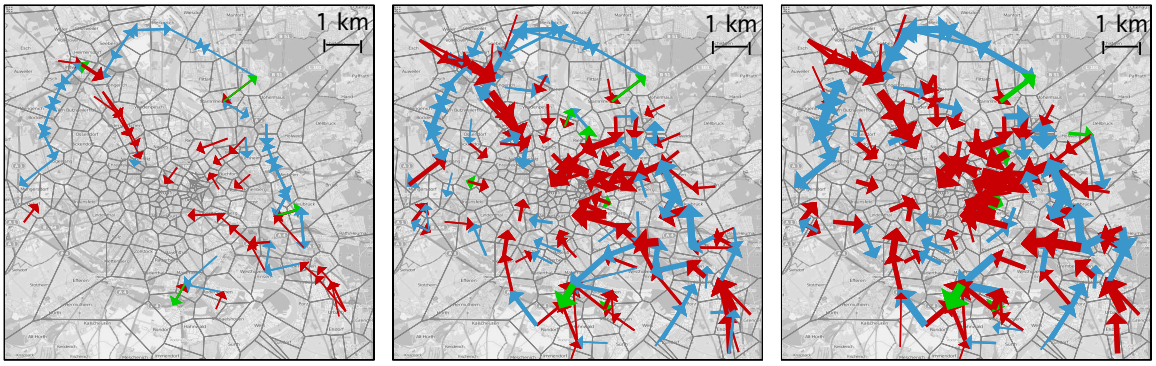

(a) 5:00 am (b) 6:00 am (c) 7:00 am

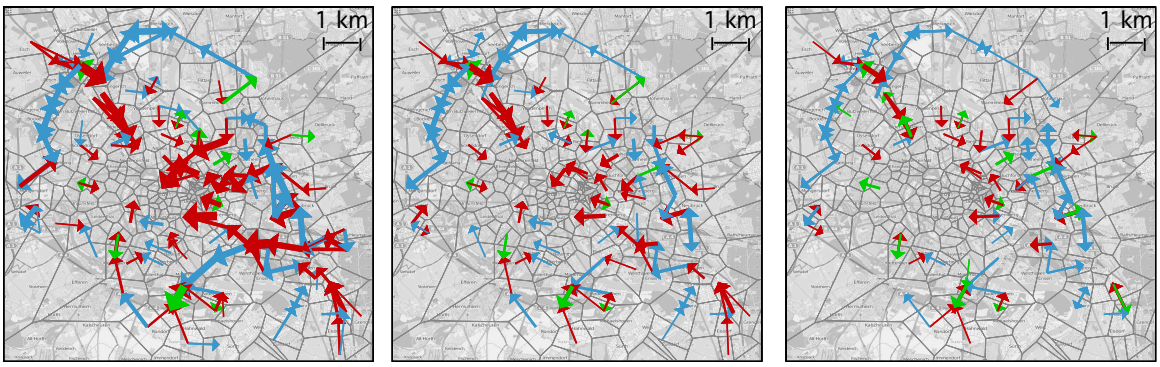

(d) 8:00 am (e) 9:00 am (f) 10:00 am

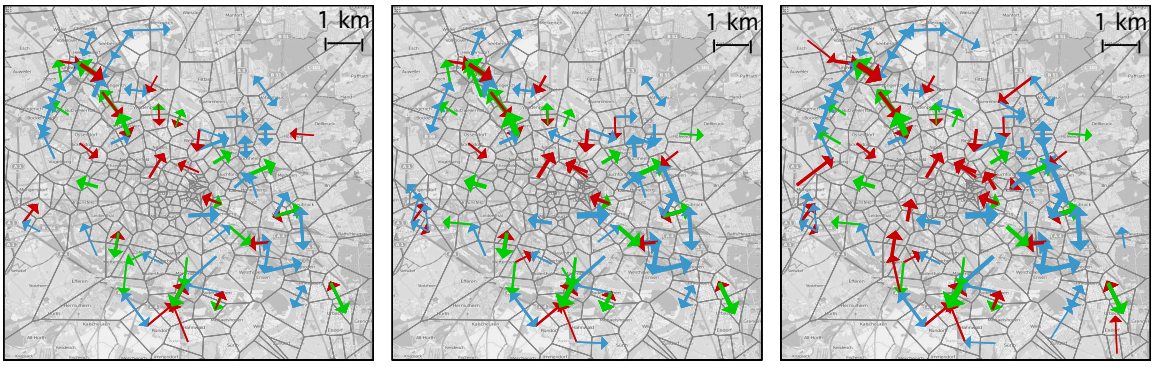

(g) 11:00 am (h) 12:00 am (i) 1:00 pm

Figure 4.5: Spatiotemporal evolution of macroscopic flows of pervasive vehicular access in the Koln region over one day. This figure is best viewed in colors (5 am to 1 pm).

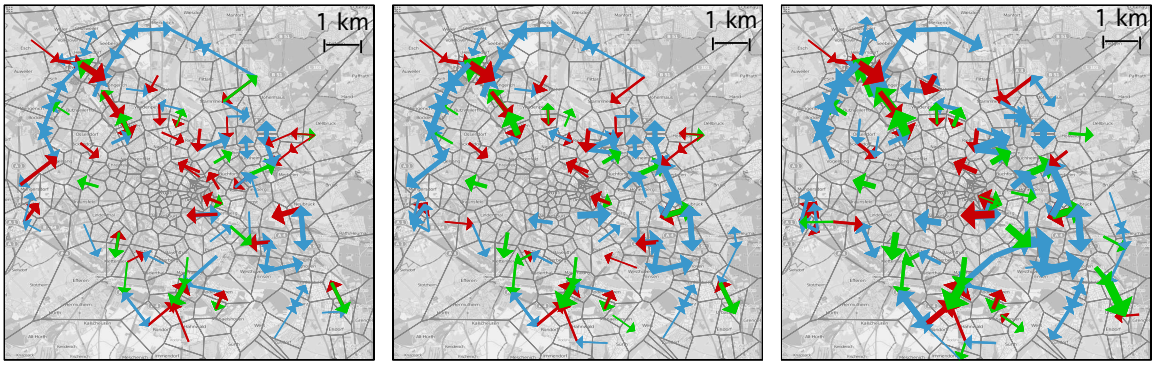

(j) 2:00 pm (k) 3:00 pm (l) 4:00 pm

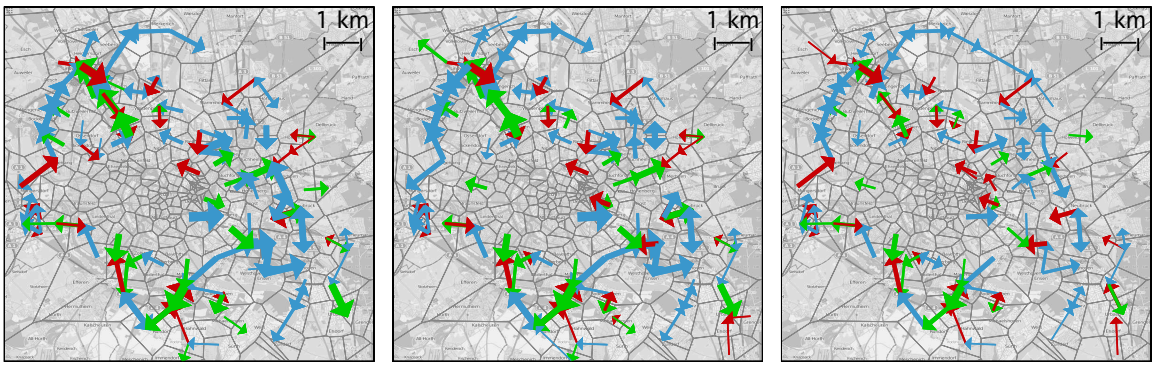

(m) 5:00 pm (n) 6:00 pm (o) 7:00 pm

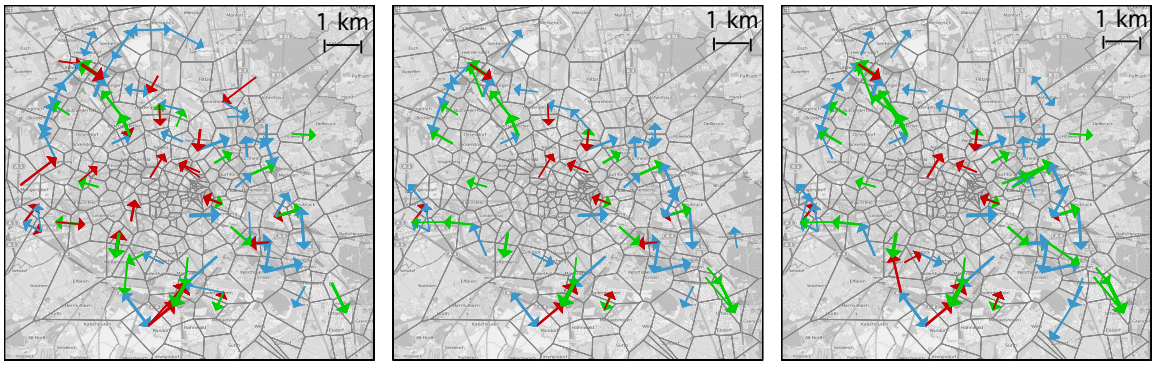

(p) 8:00 pm (q) 9:00 pm (r) 10:00 pm

Figure 4.5: Spatiotemporal evolution of macroscopic flows of pervasive vehicular access in the Koln region over one day. This figure is best viewed in colors (2 pm to 10 pm).

grow, whereas no significant outbound traffic is observable. This indicates the use of umbrella cells during these periods as the vehicular flow tend to have higher speeds, use of small cells in such cases leads to frequent handoff.

More balanced traffic headings appear during the morning, as the highway flows almost vanish, while flows within the urban agglomeration are evenly separated between the inbound and outbound directions. A similar behavior is observed until early in the afternoon.

However, from 3:00 pm to 6:00 pm, the flow dynamics change again, as the traffic volume increases during the afternoon rush hours. This time, however, the largest significant flows appear along the peripheral highways, and most of the traffic in the urban area moves away from the city center (e.g., green arrows in Fig. 4.5(m)). After 7:00 pm, traffic volumes decrease and the mobility directions become less biased and more heterogeneous.

These behaviors are consistent with the daily activity cycle in all urban areas: the commuters traffic entering the city in the morning merge with constant traffic over the highways in the morning, while it turns into flows leaving downtown in the afternoon.

From a RAN viewpoint, these results confirm the important spatiotemporal variability of the pervasive vehicular access. However, while the load density studied in Sec. 4.2 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, that feature low to moderate road traffic conditions. 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* along with designating use of different cell architectures. Indeed, a proper understanding of the dynamics of pervasive vehicular access could turn mobility into an advantage, rather than an issue to cope with.

4.2.3 Microscopic analysis

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. The properties that we consider are the vehicular user inter-arrival time, presented in Sec. 4.7,

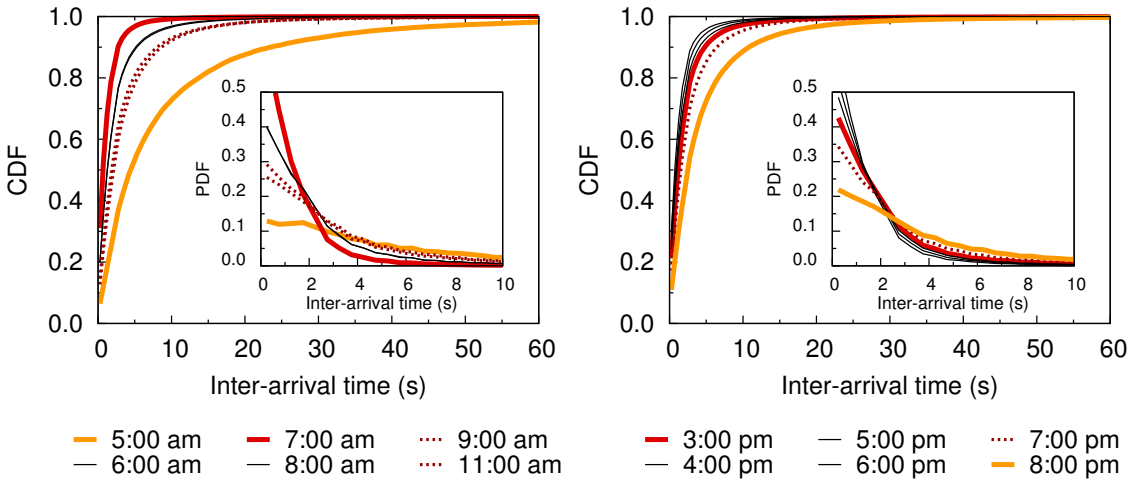

(a) Morning

(b) Afternoon

Figure 4.6: Distributions of vehicular user inter-arrival time at RAN cells.

and their cell residence time, discussed in Sec. 4.7. Both times have a major impact on the design and performance evaluation of RAN solutions for, e.g., handover management.

Inter-arrival time distributions

We first evaluate the distribution of the inter-arrival time of vehicles at each base station cell, as provided by the Voronoi tessellation. The resulting cumulative distribution functions (CDFs) are portrayed in Fig. 4.6. There, Fig. 4.6(a) and Fig. 4.6(b) resume the CDFs for the morning and afternoon hours, respectively. The inset plots show instead the corresponding probability density functions (PDFs).

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 am or 8:00 pm) result in CDFs skewed towards higher inter-arrival times, while traffic-peak hours (e.g., 7:00 am or 4 pm) result in the lowest inter-arrival times. Similarly, the CDFs overlaps for hours that yield comparable traffic conditions (e.g., from 9:00 to 11:00 am and from 4:00 to 6:00 pm).

In any case, inter-arrival times are very short during most of the day, as expected in presence of high-speed users. Between 6:00 am and 7:00 pm, at least 95% vehicular users enter the coverage area of a new cell with 10 seconds or less of each other. The time interval drops to less than 3 seconds for 50% of the vehicular users at least.

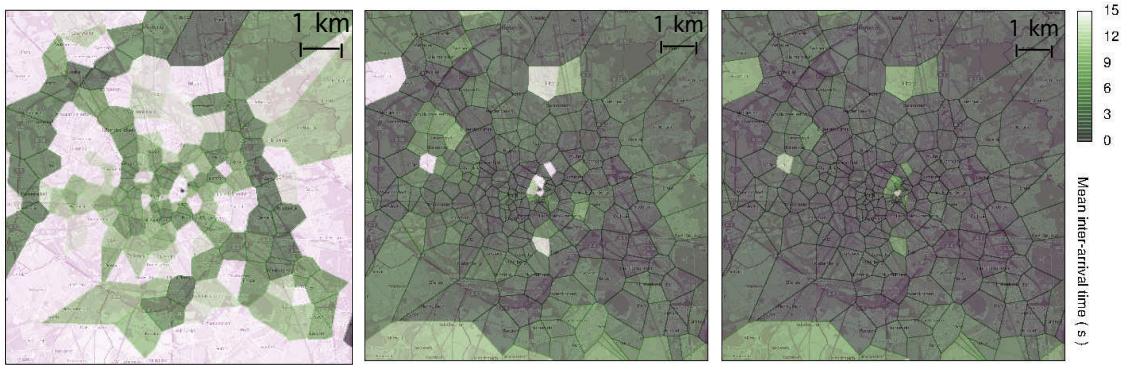

(a) 5:00 am (b) 6:00 am (c) 7:00 am

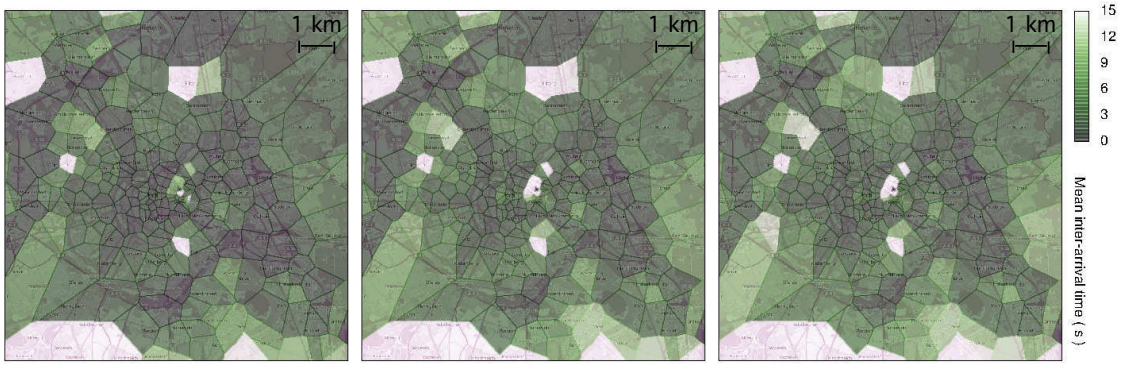

(d) 8:00 am (e) 9:00 am (f) 10:00 am

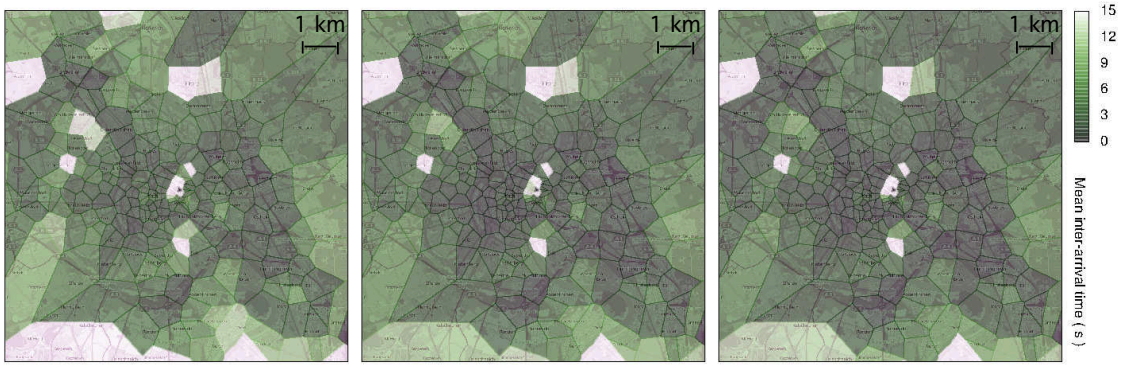

(g) 11:00 am (h) 12:00 am (i) 1:00 pm

Figure 4.7: Spatiotemporal evolution of inter-arrival times (in seconds) at cells deployed in the Koln region during over one day (5 am to 1 pm).

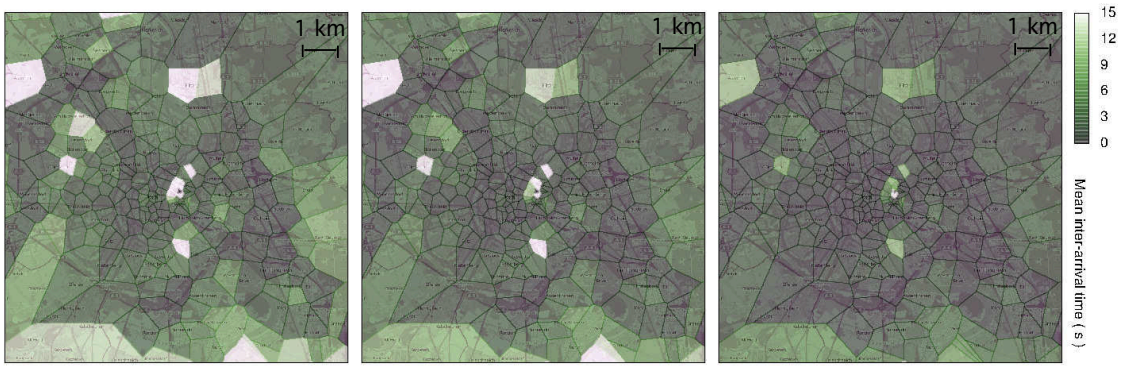

(j) 2:00 pm (k) 3:00 pm (l) 4:00 pm

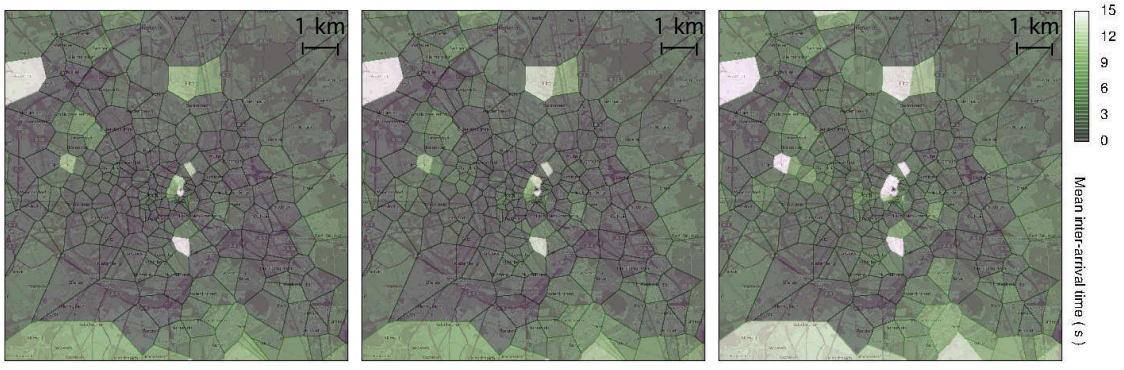

(m) 5:00 pm (n) 6:00 pm (o) 7:00 pm

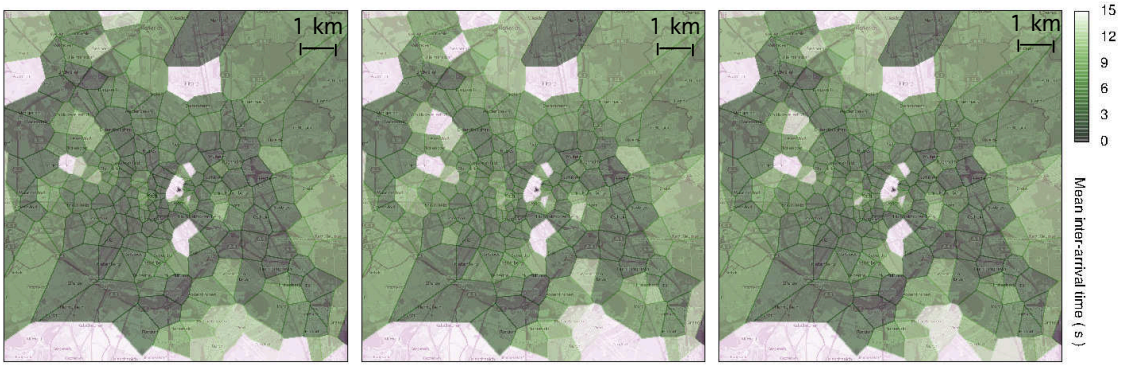

(p) 8:00 pm (q) 9:00 pm (r) 10:00 pm

Figure 4.7: Spatiotemporal evolution of inter-arrival times (in seconds) at cells deployed in the Koln region during over one day (2 pm to 10 pm).

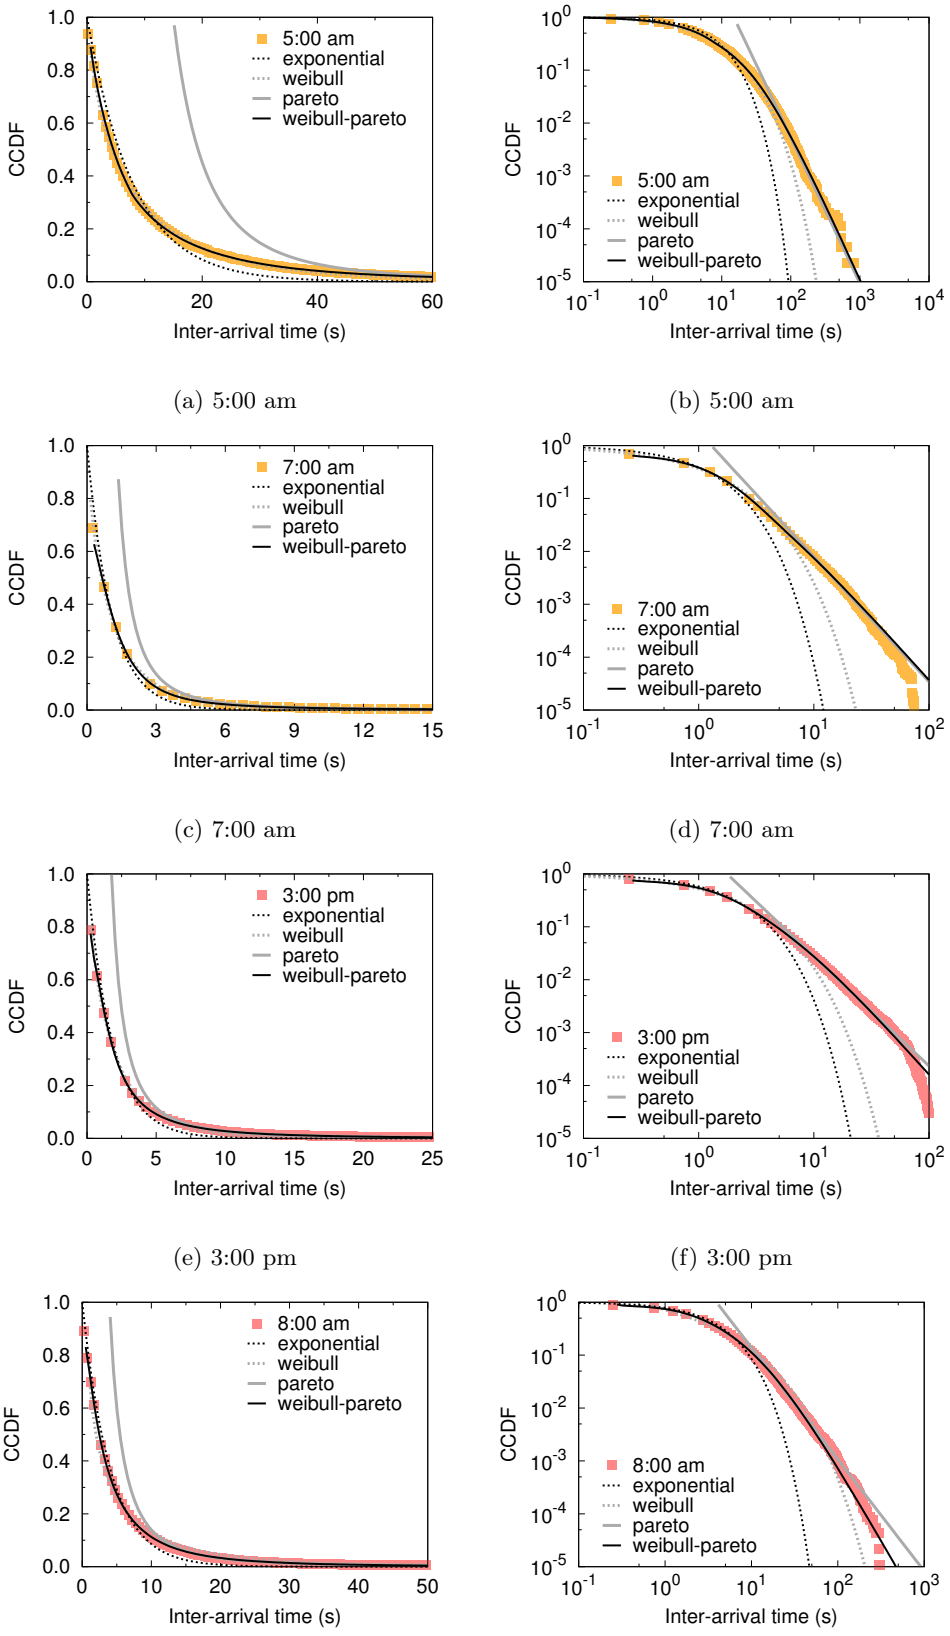

(a) 5:00 am

(b) 5:00 am

(c) 7:00 am

(d) 7:00 am

(e) 3:00 pm

(f) 3:00 pm

(g) 8:00 pm

(h) 8:00 pm

Figure 4.8: Theoretical distribution fittings on inter-arrival times in the Koln RAN deployment for different hours of a typical day. Each fitting sample is shown in linear-linear and logarithmic-logarithmic scale plots.

4.2 Access network

One could wonder how the geographical location of cells influences the inter-arrival time at each of them. Fig. 4.7 portrays the spatial distribution of average inter-arrival time values, on a per-cell basis. Variations follow the color scale on the right of each plot row, and surface when comparing different hours of the day – as we already noted from the CDFs. However, Apart from a few cells at the border of the 400 km^2 region (that lie outside the freeway ring around Koln and correspond to rural areas), we can note that the average inter-arrival time recoded at each cell is rather uniform. Indeed, cells in a single plot tend to have similar color shades. We conclude that it is mostly the temporal dimension that impacts the inter-arrival time variability, and not the spatial one.

Another interesting observation is that, although they have very diverse realizations, the CDFs and PDFs in Fig. 4.6 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 Koln scenario. A representative sample of our results is provided in Fig. 4.8, for four day times that yield different road traffic conditions: 5:00 am (very sparse traffic), 7:00 am (morning traffic peak), 3:00 pm (moderate road traffic), and 8:00 pm (mild road traffic). For each such hour, two plots show fittings in linear-linear and logarithmic-logarithmic scales: the former allow to appreciate the probability mass region, while the second let us better observe the tail of the distribution. In each plot, we compare the experimental CCDF against three well-known theoretical distributions – exponential, Weibull and Pareto – as well as one hybrid Weibull-Pareto distribution. Fittings were performed using the nonlinear least-squares (NLLS) Marquardt-Levenberg algorithm [Lev44, Mar63] .

The results confirm our intuition that the different empirical distribution follow similar laws, although with different parameters depending on the day time considered. In all cases, linear-linear plots show that a Weibull distribution fits very well the CCDFs for low inter-arrival times – noticeably better than an exponential one does. The tails of the empirical CCDFs are however heavy ones, as proven by their linear shape in logarithmic-logarithmic plots. Indeed, the Pareto class is definitely the matching theoretical distribution, and, again, an exponential distribution largely fails to reproduce the experimental behavior. Given these results, we employed a hybrid Weibull-Pareto distribution [SBrA05] to obtain a complete fitting over the whole inter-arrival time range. The exponential hybrid Weibull-Pareto distribution is defined as

$$H_{\alpha,\beta,\gamma,\delta,\rho} = H_{x \mapsto \rho^x}(X_{\alpha,\beta}, Y_{\gamma,\delta}) \tag{4.1}$$

where $X_{\alpha,\beta}$ and $Y_{\gamma,\delta}$ are random variables distributed according to standard Weibull and Pareto models given as

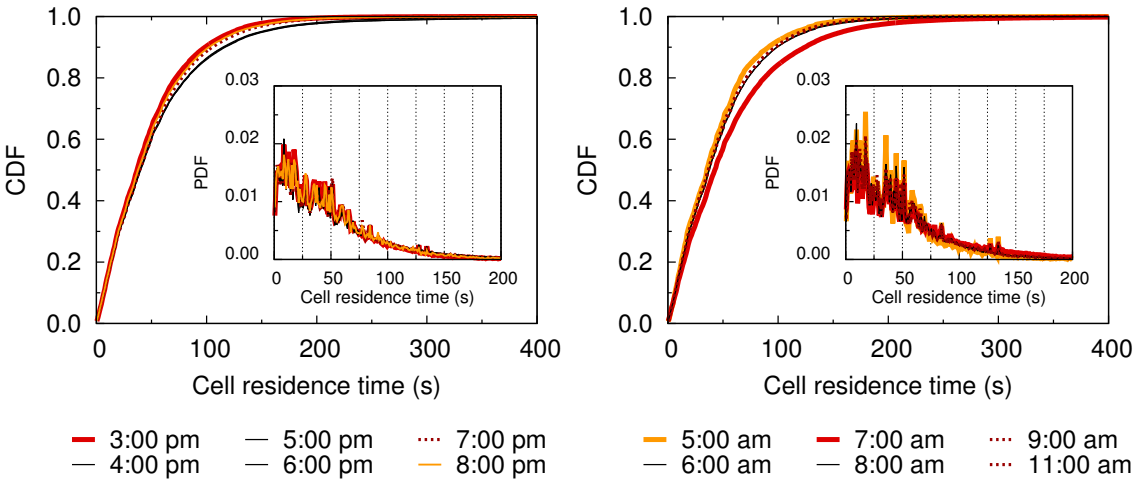

(a) Morning

(b) Afternoon

Figure 4.9: Distributions of vehicular user residence time at RAN cells.

$$F_{Y_{\alpha,\beta}}(x) = X(Y_{\alpha,\beta} \leq x) = 1 - e^{-\left(\frac{x}{\beta}\right)^{\alpha}} \tag{4.2}$$

$$F_{X_{\gamma,\delta}}(x) = X(X_{\gamma,\delta} \leq x) = 1 - \left(\frac{\delta}{\delta + x}\right)^{\gamma} \tag{4.3}$$

x in Eq. (4.1) is the weighted average of CCDF of X and Y and $x \mapsto \rho^x$ is a function that maps x to ρ^x. The hybrid distribution perfectly fits the empirical data, 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 verify the general validity of the result, they are not currently possible, due to the lack of a vehicular mobility dataset featuring a scale and a level of realism and detail comparable to the Koln one. However, our results are 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 *day time induces significant differences in the inter-arrival time distribution.* However, pervasive vehicular access results in typical *average inter-arrivals in the order of a few seconds that are geographically uniform.*

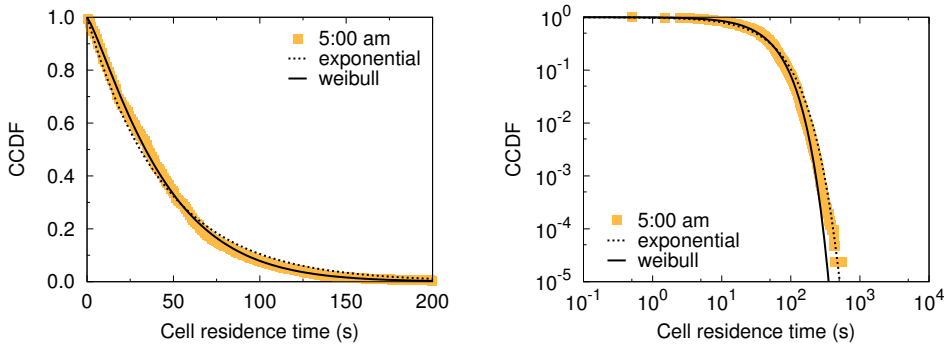

(a) 5:00 am

(b) 5:00 am

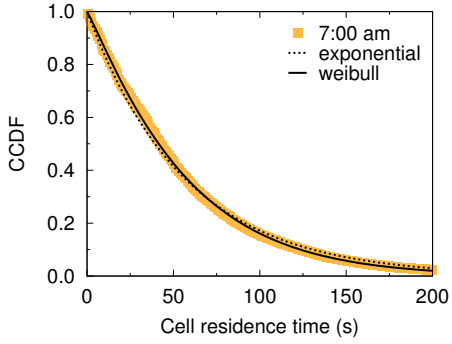

(c) 7:00 am

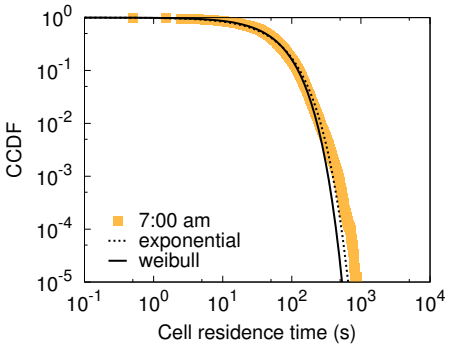

(d) 7:00 am

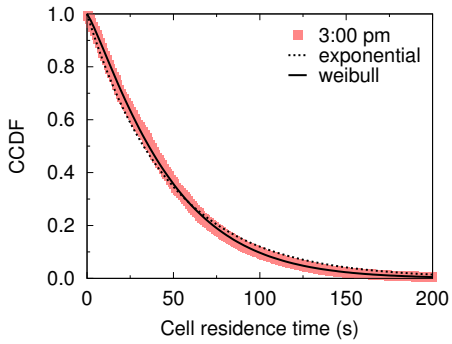

(e) 3:00 pm

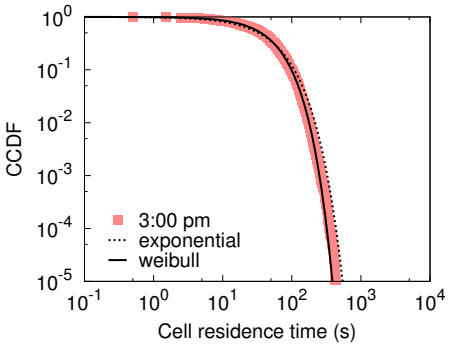

(f) 3:00 pm

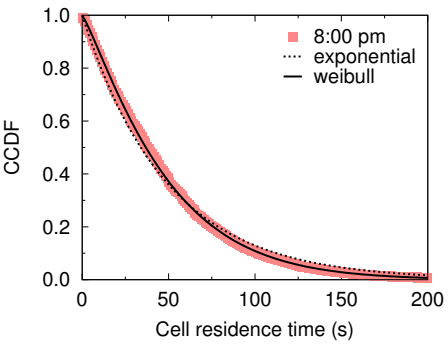

(g) 8:00 pm

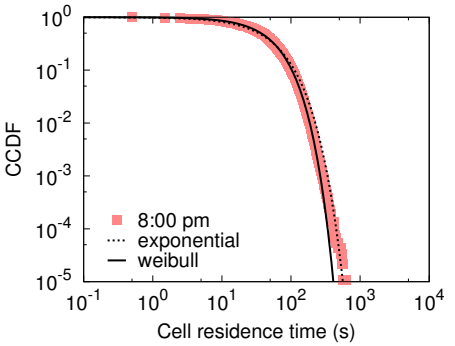

(h) 8:00 pm

Figure 4.10: Theoretical distribution fittings of cell residence times in the Koln cellular deployment for different hours of a typical day. Each fitting sample is shown in linear-linear and logarithmic-logarithmic plots.

Cell residence time distributions

Another metric commonly regarded for the dimensioning of RANs is the cell residence time [KZ97], i.e., the amount of time spent by a vehicular user under coverage of a same base station. We record the cell residence time of vehicular users in the Koln scenario, distinguishing different hours of the day. The resulting distributions are reproduced in Fig. 4.9; the two plots group CDFs and PDFs (in the inset images) for morning and afternoon hours.

Unlike what observed in the case of inter-arrival time, the cell residence time distributions measured at different hours basically overlap. In other words, the day time does not seem to have a significant impact on the time spent by vehicular users in RAN cells. In turn, this implies that, even if more vehicles arrive at a cell during the most intense traffic hours, the time spent in the cell does not vary noticeably. Although at first counter-intuitive, such a result is just an indicator that the road infrastructure can accommodate intense traffic without a dramatic impact on travel times. As a matter of fact, during the traffic peak hours, drivers tend to choose alternate routes, exploiting minor roads and spreading the traffic over the whole street layout. In other words, by better leveraging the road capacity, drivers avoid increasing their travel times – and thus the cell residence times.

In terms of absolute values, 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 leaves 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 a clear result of the high mean speed of vehicular users.

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. 4.10 displays a representative sample of the theoretical distribution fittings. We consider the same four day times considered for the case of inter-arrival times, for the reasons already discussed before. Similarly, for each hour, both linear-linear and logarithmic-logarithmic scale plots are shown.

We consider two candidate theoretical distributions, i.e., exponential and Weibull, and use the nonlinear least-squares (NLLS) Marquardt-Levenberg algorithm to best fit them to the experimental CCDFs. The outcome is that, in all cases, the Weibull distribution is that providing the best fitting. However, the exponent of the Weibull distribution is found to be very close to one, meaning that the Weibull distribution tends toward an exponential one. It is thus unsurprising that the exponential curve also fits quite well the empirical data in the plots of Fig. 4.10.

These results highlight how the common *exponential assumption may hold for cell residence times* generated by vehicular access. Moreover, the cell residence time tends not to vary significantly over the whole day, thus a single distribution instance may be sufficient to describe the overall cell residence time behavior. 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. 4.2, indicate that an important challenge lies at the intersection of the RAN capacity and the handoff frequency it induces.

4.3 Autonomous network

Having introduced a new large-scale vehicular mobility dataset in Chapter 3, that deemed to yield a higher level of realism than currently available datasets, we are now interested in understanding which impact the additional realism brought by our dataset has on the evaluation of vehicular network solutions.

In that regard, we consider few freely available datasets discussed in Sec. 2.5 namely Turin downtown, Zurich downtown, Zurich region as our reference scenarios. Additionally, to have a fair evaluation of the datasets, we try to represent our Koln trace to have similar characteristics with those of the reference scenarios. Such process involves considering Koln downtown area to match the Turin and Zurich downtown datasets, considering major arteries with less detailed microscopic simulation to match the Zurich region dataset.

In this section, in order to characterize the vehicular communication capabilities, we assume a simple disc model, with a range of 100 m, this knowledge is acquired from field tests carried out for a reliable and robust communication platform in DSRC vehicle-to-vehicle communication [BSK10]. Similar experimental evaluations in [CHS+07] reassures this considered value. This disc model allows us to point out the impact of mobility, which is our primary objective, and avoid biases due to the irregular signal propagation of urban environments.

4.3.1 Modified Koln dataset

First, we introduce the *Koln pruned* dataset, which is a simplified version of our TAPAS-Cologne dataset. It is generated using the Koln OSM road topology presented in Sec. 3.2.1, the TAPAS travel demand of Sec. 3.2.3 and Gawron's traffic assignment introduced in Sec. 3.2.4. However, it only features major road arteries as show in Fig. 4.11(a). The dataset adopts a simplified microscopic model, namely the Constant Speed Motion with Pauses model [HFFB11] calibrated so that the evolution of the average speed over time is identical to that recorded in the TAPASCologne dataset. Finally, the travel demand is limited to the morning and

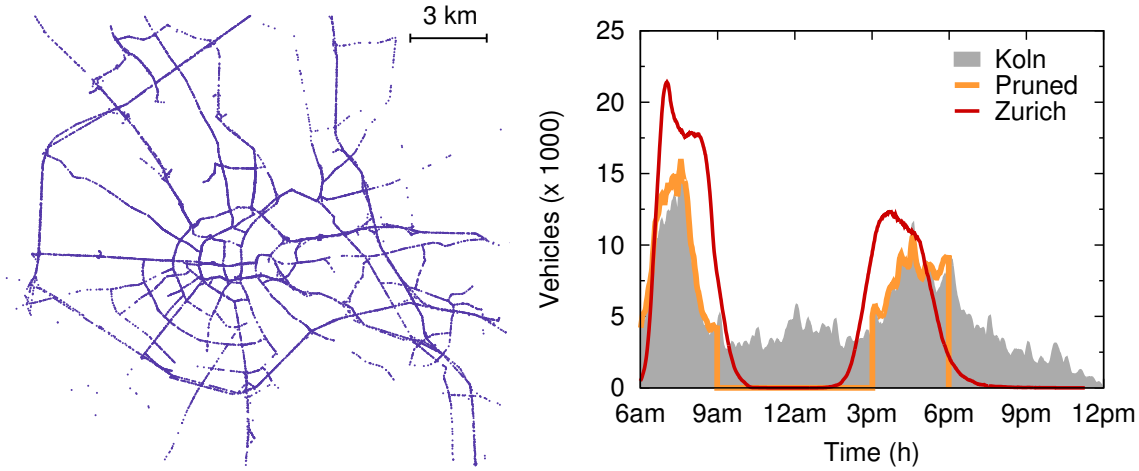

(a) Traffic snapshot

(b) Traffic volume

Figure 4.11: Koln pruned scenario: city and surroundings.

afternoon traffic peaks. With such a setting, Koln pruned dataset matches the characteristics of Zurich region (i.e., real-world road layout limited to major traffic arteries, simplistic representation of the microscopic car movement, realistic macroscopic traffic flows limited to the rush hours) as displayed in Fig. 4.11(b).

The second modified dataset is termed as *Koln downtown*. This version of the dataset captures the vehicular traffic in the inner 20 km^2 of the whole Koln region described in the TAPASCologne dataset. This dataset covers the vehicular mobility in smaller road arteries around the city centre as shown in Fig. 4.12(a).This subset of our dataset, introduced as a fair comparison term to the smaller reference datasets, Zurich downtown and Turin downtown. Traffic volume compared against these datasets is portrayed in Fig. 4.12(b).

4.3.2 Vehicular contact

Vehicular contact duration is of particular interest to the design of ad hoc or opportunistic network solutions, that assume vehicles to be able to exchange data directly, i.e., without passing through a roadside network infrastructure. Clearly, the duration of contacts among vehicles plays a major role in determining the amount of data that can be exchanged among cars, and thus the potential of ad hoc or opportunistic communication.

Fig. 4.13 shows the PDF of the inter-vehicle contact time measured over the whole region during the 24 hours. We can note that contacts are typically extremely short: most of the probability mass gathers between 1 and 15 seconds in case of Koln region, whereas in Zurich

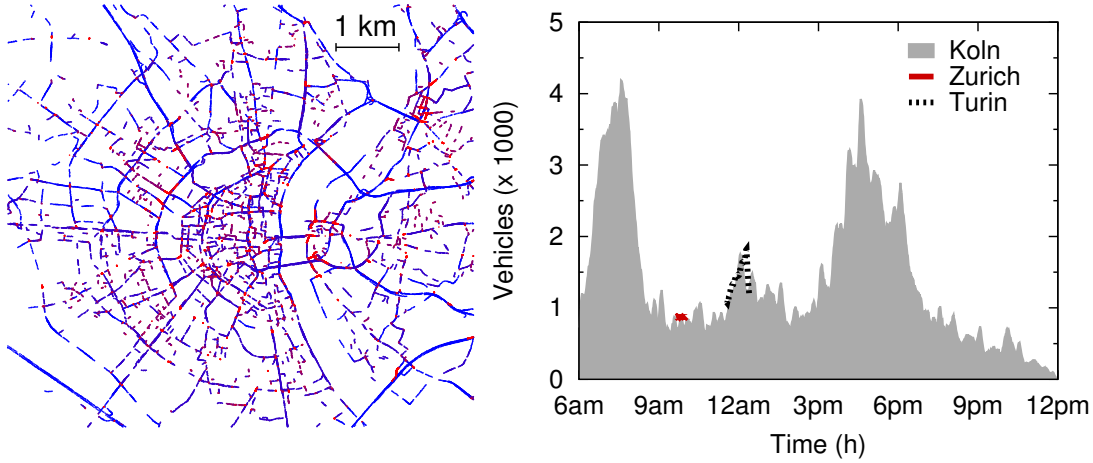

(a) Traffic snapshot

(b) Traffic volume

Figure 4.12: Koln scenario: downtown.

region most of the contacts span from 1 to 10 seconds with majority of them with 5 seconds duration. However, the distribution is heavy-tailed in Koln region. This is more clear in the inset plot, portraying the complementary CDF (CCDF), i.e., the probability that the contact lasts more than the value on the abscissa, in a logarithmic scale. From the CCDF, we can observe that, although with a much lower probability, contacts can last up to several minutes, where Koln region looks better compared with Zurich region.

These results make us conjecture that exploiting contacts among vehicles to exchange data may be a challenging task. The vast majority of contacts last a few seconds: considering the time required to identify the presence of a new neighbor and to establish a link, these contacts are hardly usable for any application other than periodic broadcasting. However, long-lasting links are also present: if properly identified, they can be leveraged for the exchange of large amounts of data.

4.3.3 Network graph metrics

Defining cluster as a group of vehicles such that a (multi-hop) path exists between any pair of them at a given time instant [FH08]. Therefore, vehicles belonging to different clusters cannot communicate at that time, neither directly nor passing through other vehicles. Fig. 4.14 shows the average number of clusters observed in each scenario. The plot also reports the mean and standard deviation of the cluster size, i.e., the number of vehicles that belong to a cluster.

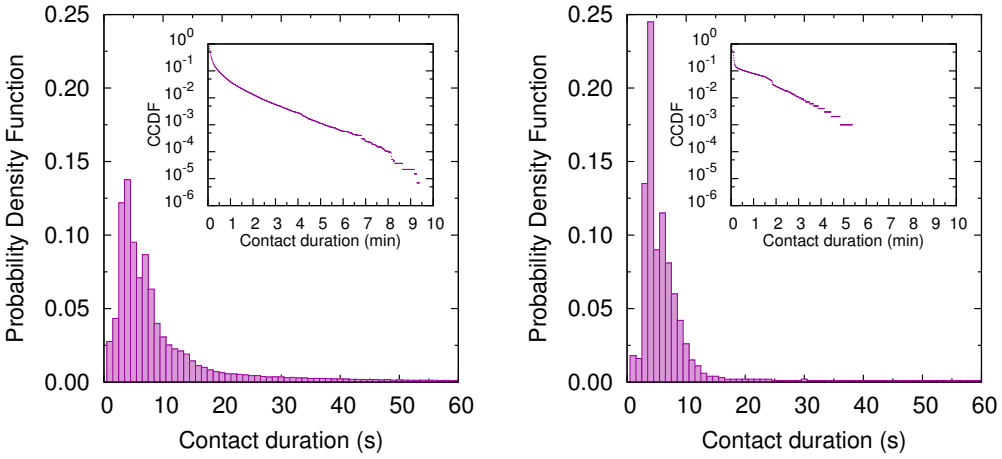

(a) Koln region

(b) Zurich region

Figure 4.13: Distributions of the inter-vehicle contact duration, over the whole day.

By looking at the larger *region* scenarios, on the left of the plot, a clear difference emerges between our dataset and that of Zurich: the latter results in a much more connected network than the former, with vehicles grouping in less than one third of the clusters we record in our dataset. One could wonder whether that is the effect of a much higher percentage of singletons, i.e., clusters composed of one isolated node, in the Koln dataset: from the figure, however, a similar fraction of singletons is present in both scenarios, accounting for approximately 60% of the overall clusters. The reason for such a difference is instead explained by the extremely high average and standard deviation of the cluster size in Zurich scenario[1]: these are evidences of the coexistence of the many singletons with several giant components that gather a large portion of the vehicles. Such giant components cannot instead be found in the Koln scenario, where clusters tend to be much smaller and more uniform in size.

Such an average value analysis is confirmed by the results in Fig. 4.15. The left figure, detailing the evolution of the cluster number over time, shows that the behavior previously described is actually not influenced by the daytime. Indeed, the number of clusters in the Zurich scenario is significantly lower than that recorded in our Koln dataset at all times. The plot also confirms that the Zurich dataset is unusable between 10:00 am and 2:00 pm, as well as after 8:00 pm.

The right image of Fig. 4.15 confirms instead our intuition on the presence of giant components in the Zurich dataset. As a matter of fact, the cumulative distribution function

[1]The cluster size standard deviation in the Zurich region and Koln pruned scenarios exceeds 290 nodes, and was not included for the sake of clarity.

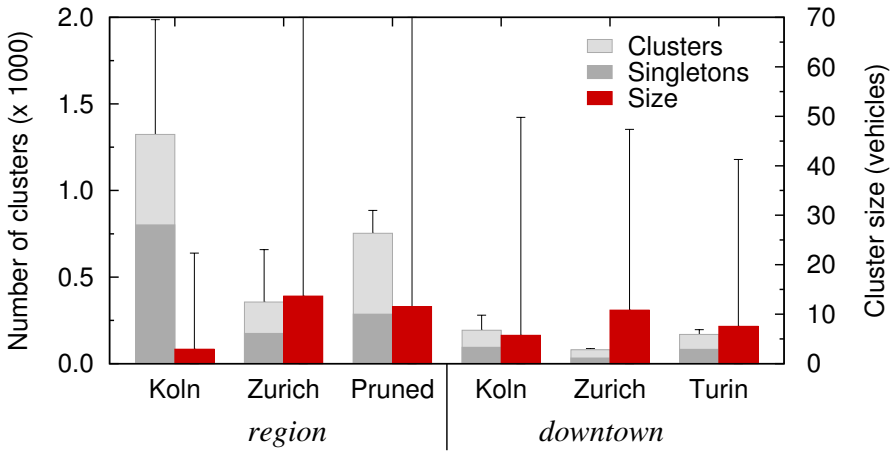

Figure 4.14: Average number of clusters, with singletons, and mean cluster size.

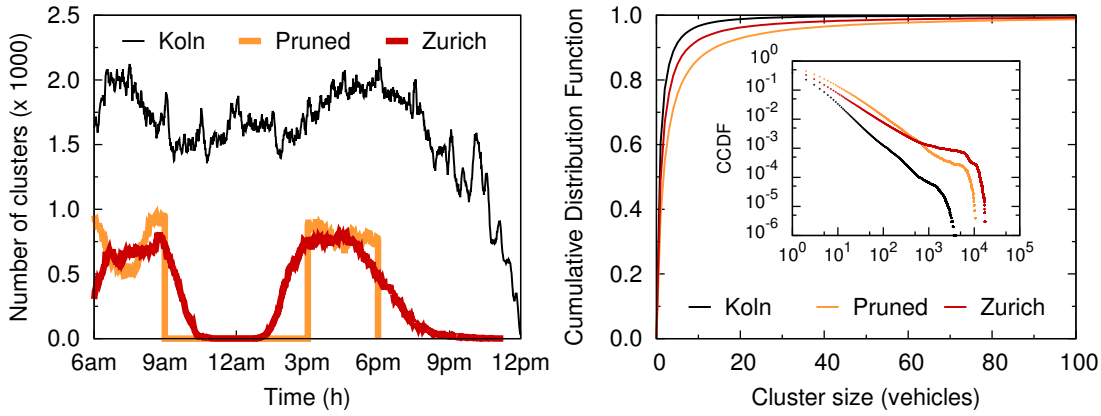

Figure 4.15: *Region* scenarios. Clusters in time (left) and cluster size CDF (right).

(CDF) of the cluster size, in the outer plot, proves that the Koln dataset contains a higher number of small (20 vehicles or less) clusters. Conversely, the complementary CDF, in the inset plot, highlights the tail of the distribution, allowing to observe a much more probable (0.1%) presence of very large clusters (more than 10000 vehicles) in the Zurich scenario. In the Koln dataset, the largest cluster does not exceeds 4000 cars, and appears with orders-of-magnitude lower probability (0.001%), even if the number of cars concurrently traveling in the two datasets is comparable, see Fig. 4.11(b).

At this point, one could rightfully ask whether the connectivity differences between the Koln and Zurich region datasets are imputable to the diverse urban areas represented in the two mobility datasets. Indeed, the two scenarios are characterized by dissimilar road topologies and road traffic flows, which could justify the non-comparable connectivity results. However, let us observe the behavior of the Koln pruned dataset in Fig. 4.14 and Fig. 4.15. Despite the fact that it portrays road traffic in the Koln area, the Koln pruned dataset results in a

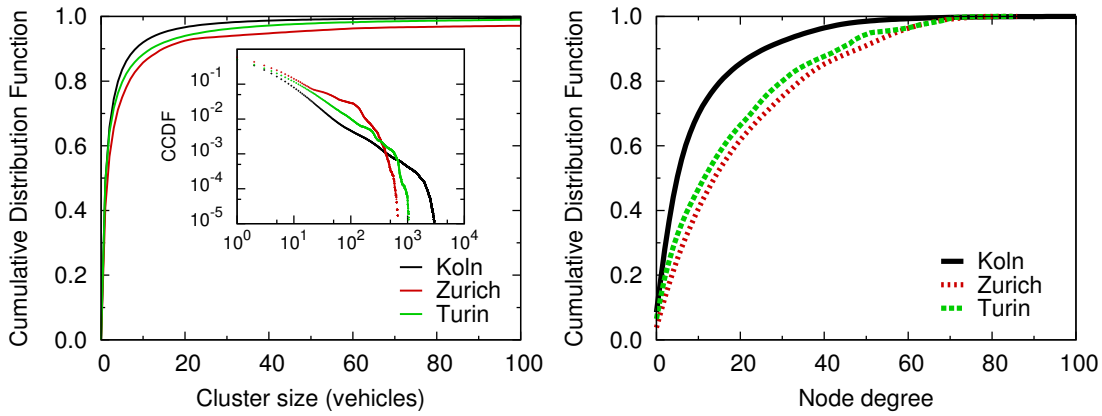

Figure 4.16: *Downtown* scenarios. Cluster size (left) and degree (right) CDF.

vehicular network featuring a time-varying number of clusters nicely matching that of the Zurich region dataset. Moreover, it yields a low average cluster size with a very high standard deviation as well as 10000-node clusters appearing with 0.1% probability, exactly as in the Zurich region case.

Therefore, we conclude that it is not the underlying urban environment that determines the differences between the Koln and Zurich region scenarios, rather the diverse level of realism of the mobility description. In particular, the emergence of unrealistically large components in the Zurich dataset is imputable to the combination of reduced road topology information and low microscopic mobility detail, the same features we observe in the Koln pruned dataset. Considering only major roads and a simplistic microscopic approach leads therefore to a very homogeneous traffic over the few traffic arteries, that act then as seamlessly connected backbones for the vehicular ad hoc network.

As far as the *downtown* scenarios are concerned, their average behavior in terms of clustering is depicted in the right portion of Fig. 4.14. The number of clusters is significantly lower than that observed in the larger *region* scenarios, consistently with the reduced size of the areas. In this case, results are more similar through the different datasets, however the connectivity of the vehicular network in the Koln dataset presents a slightly higher variability in both cluster number and size: this is due to the fact that our dataset captures the evolution of the traffic over the day, whereas the other scenarios are only representative of a short time span characterized by quasi-static network clustering properties. The left plot in Fig. 4.16 supports such a conclusion. There, we also remark that the largest component sizes observed in the Koln region and downtown scenarios match: such a component appears during the rush hours in the trafficked city center.

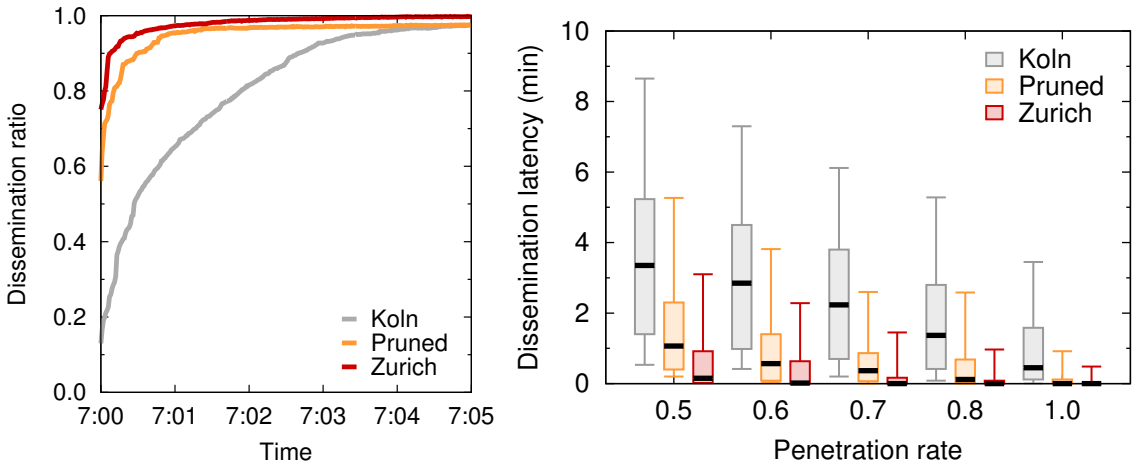

Figure 4.17: Epidemic dissemination ratio over time with a penetration rate of 1.0 (left) and latency to achieve the dissemination ratio quantiles 0.05, 0.25, 0.5, 0.75 and 0.95 under different penetration rates (right).

The right plot in Fig. 4.16 focuses on the node degree, i.e., the number of communication neighbors of a vehicle, an especially relevant metric in small-scale scenarios. We can note that vehicles in the Koln scenario tend to have smaller 1-hop neighborhoods, with only 5% of the them having more than 30 neighbors, while 60% have less than five nodes within communication range. On the contrary, in the reference downtown scenarios, the fraction of vehicles with large neighborhoods of more than 30 nodes grows to 25%, and only 20 to 30% of the nodes have five or less neighbors.

Summarizing our findings, we can conclude that the topology of a vehicular network built on the car traffic described by our dataset is sensibly different from those obtained with currently available mobility datasets. More precisely, when compared with the standard large-scale mobility dataset, i.e., Canton of Zurich, our dataset appears significantly more detailed. In turn, such additional detail leads to a less connected and less stable network. Considering instead small-scale datasets, our dataset shows an equivalent level of detail, but a much more variegate behavior, as it models a whole day rather than a few tens of minutes of road traffic. As a result, the Koln dataset allows to observe a connectivity variability that is not captured by the other datasets. Overall, the discussion brings us to conjecture that evaluating network protocols or architectures through low-detail or spatially- and temporally-limited mobility datasets is a risky practice, that can lead to over-optimistic performance results. Network simulation results in the next subsection substantiate this conclusion.

4.3.4 Epidemic dissemination

We now consider a networking application use-case, and evaluate the effect that a high-detail large-scale mobility dataset has on the system performance. More specifically, we focus on the epidemic dissemination of some small content (e.g., information on the current status of road traffic in the area or a map update for the on-board navigation system) throughout the whole network. We assume a susceptible-infected model, where a node that has been reached by the content (infected) can forward the data to the uninformed vehicles (susceptible) it encounters. We consider that the content can be transmitted over the wireless channel in a few milliseconds, and that transmissions occur in a broadcast fashion, so that all 1-hop neighbors of the sender can receive it at once. Finally, we neglect medium contention and channel errors: we recognize these to be strong assumptions, however (i) they make simulations computationally feasible in the very large-scale scenarios we consider, and (ii) they do not significantly impact our comparative evaluation, since they affect in a similar manner the different mobility datasets.

The epidemic dissemination is run in the 400 km^2 *region* scenarios, i.e., the Koln, Zurich region, and Koln pruned datasets. A source, located at the city center, 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.

The left plot of Fig. 4.17 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 Koln dataset is more gentle, with two minutes required to reach 80% of the network, the Zurich region and Koln pruned datasets lead to a much faster spreading. In fact, 80% of the vehicles are informed in a few seconds in these scenarios. This is an artifact of the unrealistically high connectivity that, as our previous topological analysis unveiled, characterizes the Zurich region and Koln pruned vehicular networks: it is indeed the presence of large clusters including most of the road traffic that makes the spreading exceedingly fast.

The right plot of Fig. 4.17 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 the *region* scenarios. The result shows that the content is successfully disseminated throughout the whole network in all cases, however it is also clear that the very high connectivity granted by the Zurich dataset leads to latencies that, under

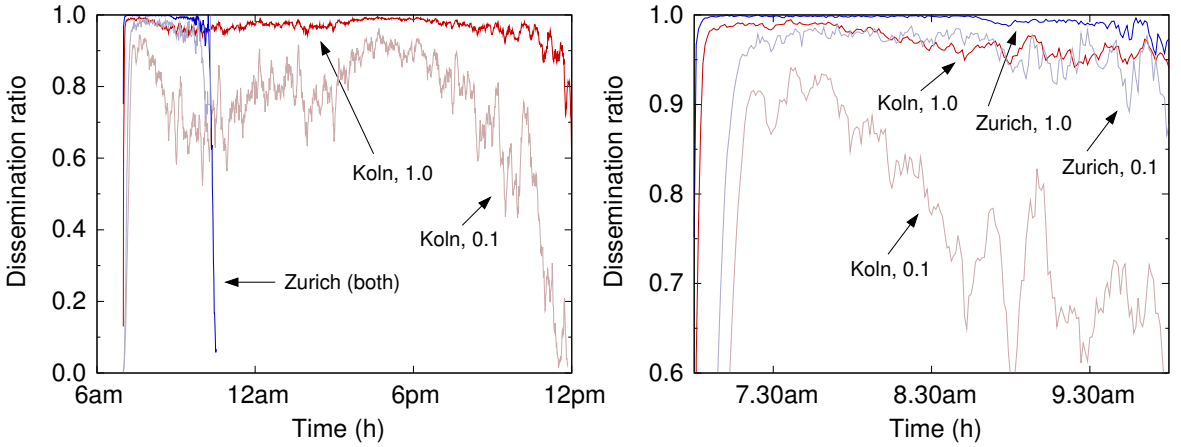

Figure 4.18: Survivability of the epidemic dissemination. Penetration rate over the whole day (left) and during the morning rush hours (right)

the different penetration rates, are from two to six times lower than those recorded in the Koln scenario. Interestingly, the same is true for the Koln pruned dataset: this confirms that it is the realism of the mobility description that leads to the network performance difference, rather than the specificities of the urban areas considered.

Not only the latency, but also the content survivability is a metric of interest in the study of epidemic dissemination. The latter represents the capacity of the content to self-sustain in the network, so that new vehicles starting their trips can immediately receive it. The left plot of Fig. 4.18 portrays the evolution of the dissemination ratio over the whole day, after the initial injection at 7:00 am. The curves refer to the Koln and Zurich datasets, since the Koln pruned dataset performs once more close to the Zurich one and is omitted for the sake of clarity. Penetration rates of 1.0 and 0.1 are considered. The figure clearly shows how the content dies out at around 10:00 am in the Zurich scenario, whereas it is able to self-sustain during the whole day in the Koln dataset. Indeed, the sudden drop in the dissemination ratio observed in the first scenario is imputable to the disappearance of road traffic after the morning peak: the complete absence of vehicles after 10:00 am makes it impossible for the content to survive.

However, as long as traffic is present in the Zurich network, the elevate connectivity leads to extremely good survivability performance under any penetration rate. In the right plot of Fig. 4.18, we provide a zoom on the morning rush hours, that outlines how the dissemination rate is constant at 100% in the Zurich dataset, when all vehicles participate in the system. In fact, such performance is only slightly affected by the penetration rate in the Zurich scenario, as limiting the network to 10% of the nodes still allows to maintain the content alive throughout 95% of the network. It is interesting to note that the dissemination self-sustains

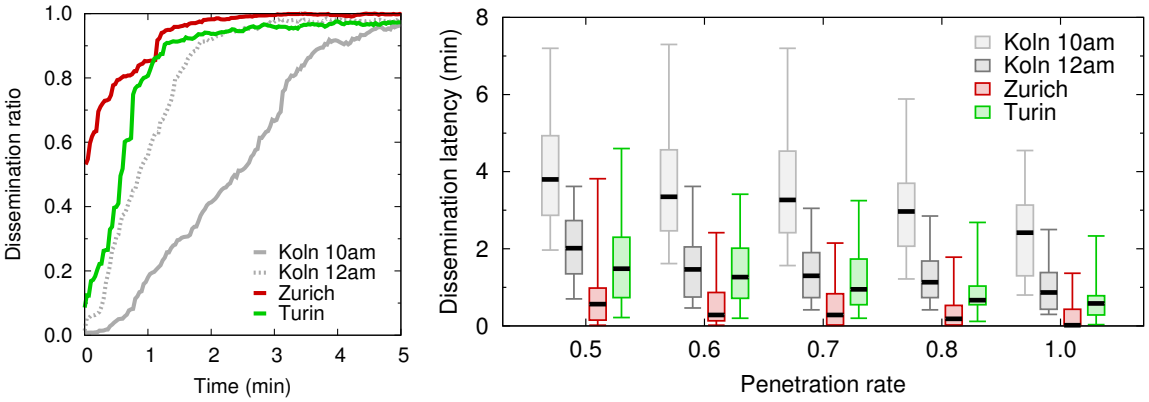

Figure 4.19: Initial spreading of the epidemic dissemination. Evolution of the dissemination ratio over time when the penetration rate is 1.0 (left) and latency to achieve the dissemination ratio quantiles 0.05, 0.25, 0.5, 0.75 and 0.95 in presence of different penetration rates (right).

in a comparable manner in the Zurich scenario with 10% of the nodes and in the Koln one with 100% of the nodes. In the latter scenario, reducing the penetration rate to 0.1 results instead in a dramatic loss of performance, the content surviving in 60% of the network only.

These results confirm that limited realism in the representation of the vehicular mobility can have a significant impact on the performance evaluation of networking solutions.

We now compare the performance of the epidemic dissemination protocol in the different *downtown* mobility datasets with two hours of the day, namely 10:00 am and 12:00 am. The left plot of Fig. 4.19 portrays the content diffusion process in presence of a penetration rate of 1.0. The plot shows the ratio of vehicles reached by the content over time. We can observe that the Zurich and Turin datasets lead to faster dissemination of the information than the Koln dataset, at both 10:00 am and 12:00 am. However, while the difference is remarkably large between the Zurich scenario and its Koln equivalent (in terms of traffic volume) at 10:00 am, the same is not true in the case of Turin. In fact, the epidemic process in the Turin dataset is only slightly faster than that in the Koln region. This is due to the fact that the Turin dataset was generated using OSM data and the SUMO microscopic mobility simulator, exactly as our Koln dataset: therefore the two scenarios feature the same small-scale realism, and differences lies in the description of macroscopic traffic flows. However, the impact of realism in the travel demand and traffic assignment is not always easy to appreciate over small geographical areas such as those considered in the Turin and Koln downtown datasets. This is why we consider the epidemics performance to be more significative on the large-scale *region* scenarios that in the small-scale *downtown* ones.

Theses observations are confirmed when studying the behavior of the epidemics at different technology penetration rates, in the right plot of Fig. 4.19. There, we can note again the large difference between the time needed to inform the 5%, 25%, 50%, 75% and 95% of the vehicles in the Zurich and Koln scenario at 10:00 am. Such a difference is instead less significant between the Turin dataset and the Koln dataset at 12:00 am.

Finally, we remark that the evaluation of information survivability as in Fig. 4.18 is of no interest in the *downtown* scenarios, since both the Zurich and Turin datasets are too short for the discussion to be meaningful. In other words, information self-sustain in such scenarios until the mobility description abruptly ends a few tens of minutes after its start.

4.4 Summary

In this Chapter, we present an analysis of the effect of vehicular mobility dynamics on the cellular network setup that is realized with the help of real cellular base station deployment and Voronoi tessellation. The macroscopic analysis unveiled that the pervasive vehicular access has a strong variability with the time of the day that cannot be simply derived by mapping the population density or mobile device distribution. Also we identified consistent vehicular traffic flows between cells that last for most part of the day, and this results can be used efficiently to configure resource allocation by spatial analysis. The microscopic analysis shows that the inter-arrival times does not necessarily follow a Poisson distribution as commonly assumed in most of the current research in cellular technology. Due to the high speed of the vehicular users, low cell residence times are observed with peak hours of the day, this hints at the use of umbrella cells in certain geographic areas of the city so as to decrease the number of frequent call handoff.

In the second part of this Chapter, we presented the analysis of the effect of road traffic dynamics from the autonomous network point of view. By reducing the realism of our dataset to match several reference scenarios, we demonstrated how realism in traffic dynamics must be considered in wireless networking. This study also highlighted that our realistic scenario is often less connected and less stable hinting at the need for V2I communication architectures especially during specific times of the day in certain geographic areas of the city for time critical applications.

Online macroscopic vehicular mobility prediction

5

Contents

5.1 Introduction 90
5.2 Related work 91
5.2.1 Basic prediction approach 91
5.2.2 Vehicular mobility prediction on road networks 92
5.3 Macroscopic vehicular mobility prediction models 94
5.3.1 Markovian model 94
5.3.2 Sequence Matching model 98
5.3.3 Online prediction with fallback 99
5.3.4 Estimating the prediction accuracy 100
5.4 Scenarios 102
5.4.1 Vehicular mobility adaptation for prediction model. 103
5.4.2 Impact of model parameters for crossroads 104
5.4.3 Impact of model parameters for the cellular network scenario 107
5.5 Performance evaluation 110
5.5.1 Road network scenario 110
5.5.2 Cellular network 120
5.6 Summary 126

5.1 Introduction

The increase of population in urban areas over the last decade has shifted the role of proactive communication and transportation system management from an option to a requirement. Anticipating the user demand in wireless access networks, or traffic congestion in road networks has emerged as an important topic of research. For instance, when a cellular user moves towards the edge of the coverage area of his current cell, handoff occurs to a neighboring cell; if there is not enough capacity, the call or media streaming is dropped or interrupted. It is possible to reserve resources at all neighboring cells for handoff users: however, this scheme leads to an overall waste of resources, and makes new calls suffer from a severe blocking problem even at light loads. Hence, pre-configuring the resources by anticipating user mobility plays an important role in providing uninterrupted reliable services. In an autonomous vehicular network scenario, future location prediction information can aid in maintaing a robust network or increasing the reliability of delivery of information in multi-hop communication.

In this Chapter, we take advantage of the Koln dataset to evaluate the effectiveness of vehicular mobility prediction in presence of realistic road traffic. Initially in Sec. 5.2, we present several closely related works on vehicular mobility prediction and discuss their shortcomings. In Sec. 5.3, we discuss our model, and the parameters that we use to calibrate it. We derive motivation from the dynamics of macroscopic vehicular flow patterns presented in Sec. 4.2, to propose a model that predicts future locations based upon the knowledge acquired within a small window of historical data. Then, we discuss the significance of the parameters that are essential for a mobility prediction model and present the calibration of the same. In Sec. 5.4, we study the predictability of vehicular mobility between crossroads in a road network and between cells in a cellular network. Performing the prediction at crossroads makes this work general, as results can be applied to RAN architectures, autonomous networks, intelligent traffic management systems, and many more. Predicting vehicular mobility across cells is more concentrated towards cellular networks as we employ real base station deployment for our study. Both these scenarios provide an understanding of how calibration and analysis of the prediction model parameters are significant while designing a prediction system for vehicular mobility. In Sec. 5.5, we discuss the spatiotemporal properties of the mobility prediction over the geographic area of Koln. The results highlight the direct relation of prediction accuracy with the volume of mobility transitions at a network node. Our results prove that shorter historical knowledge can predict the immediate future with similar accuracy as obtained with large historical knowledge in past research attempts.

5.2 Related work

Mobility prediction has long been a topic of research in various research communities, exposing new ways of extracting human mobility information. Algorithms from various disciplines have been proposed to achieve a good prediction accuracy with low processing and time complexity. In this Section, we highlight the basic approach involved in mobility prediction and review several closely related works in vehicular mobility prediction. In addition, we discuss the relevant properties that need to be considered while predicting vehicular flows.

5.2.1 Basic prediction approach

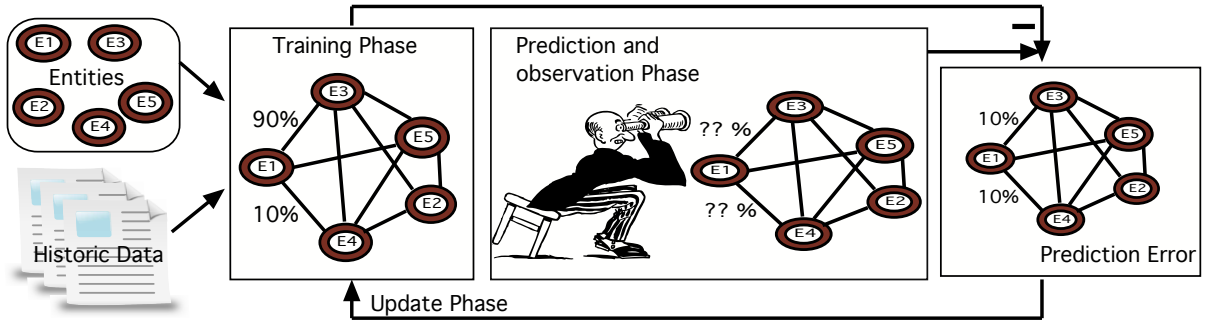

Figure 5.1: Different phases in basic prediction process.

A basic prediction system can be realized as shown in Fig. 5.1. Initially, the system is provided with historical data and the entities to be learned. Now, the system operates through three main phases. The *training phase* involves building a relational model between the provided entities based on a prediction algorithm. Once this training phase is completed, the system is ready to predict the future possible actions or states with reference to the relational model and the current situation. In parallel to prediction, the system also observes the current situation to monitor the correctness of the system which forms the *prediction and observation phase*. In case of differences in the predicted value, the difference is fed to the prediction algorithm that learns about this behavior and updates the relational model in *update phase*.

In human mobility prediction, the prediction system is typically built on entities such as unique visited locations as entities like home, office, and various point-of-interests. Also additional information like time-of-day, age of the person, mode of transport and many more details are employed to add precision in prediction depending upon the context. To derive the relational model between the aforementioned entities, many algorithms have been experimented in early works including techniques such as Data compression [VK96, BD99, GC07], Data min-

ing [YKUM05, JLSZ08, JYZJ10, WLY+11], Neural networks [VKV02, CB04, AS07, Civ06], Markovian chains [AS02, SKJH04a, DK04, Las06, SDK+06, AMSS11, GKdPC10], Sequence matching [LM96, CZS98, ESE+01, PKSG11], and many more. The vehicular mobility prediction that we focus on can be intended as a subset of the human mobility prediction techniques above and are introduced next.

5.2.2 Vehicular mobility prediction on road networks

Many attempts have been made to capitalize on the constrained mobility of high-speed cellular users for intelligent use of available resources to provide better quality of service. Previous vehicular mobility prediction attempts that are close to our work are discussed next highlighting the difference and the significant properties that need to be considered during prediction.

In [XLYL12], the vehicular mobility dataset of 4,000 taxi's in Shanghai region is used to extract the Vehicular Mobility Pattern (VMP) along road segments. Later, transition probabilities between these patterns are calculated and are used when selecting an ideal candidate to forward a packet in vehicular autonomous network. This work includes real GPS datasets from taxis, which attract map matching algorithm whenever the probabilistic model is generated or updated, hence adding to processing complexity. Moreover, taxis often follow regular routes, there by limiting the mobility prediction to fewer road segments. Our analysis is more extensive with more than 27,615 crossroads covering 400 km^2 of Koln metropolitan region including around 4500 km of roads. Hence it provides a complete description of diverse macroscopic flows.

In [PLFK03], an attempt is made to capture and predict the high-level behavior of individuals, like the mode of transportation and the route followed in an urban environment. The model uses a particle filters to build the mobility patterns of a user with bits of travel information performed through different modes of transportation. Any irregularities in the user movements are addressed by random variables. An Expectation-Maximization technique is adopted to build a relational model between modes of transportation. The training data are drawn from a GPS sensor stream collected by the authors over three months. To evaluate the prediction, they present the prediction accuracy in terms of prediction length in number of city blocks. The model accuracy descends from 100% to 52% when predicting from one to five city blocks ahead from a given block. Since individual mobility information is studied, the objective of predicting the transportation mode followed by individuals and scalability of the approach remains questionable. Our approach closely relates to this work, but we concentrate on predicting groups of commutes performed by individuals at a larger aggregate scale.

5.2 Related work

In [Sch07], hexagonal cells with the diameter of 200 meters were used to cover the geographic area of Aalborg, Denmark, and a Bayesian network was constructed with these cells. Predicting the mobility of 20 different users provides 80.54% prediction accuracy. In another similar study, [Kru08] performs the prediction at crossroads with mobility data from 100 drivers over 12 days. This work shows that higher order Markovian models, as high as *order-10*, to perform the best with 90% prediction accuracy considering substantial decrease in predictable cases. There, the prediction is confined to individual drivers, hence raising the privacy issues. Both these studies involve a small number of individuals, although their movement patterns are collected over several weeks or months. As people often follow some kind of regularity in their mobility [LM96], approaches with such data are more predictable, whereas the Koln mobility dataset includes the mobility patterns of a city with more than 900,000 inhabitants: one can imagine the degree of diversity that such a dataset yields.

In the light of all these attempts, we argue that for a detailed useful analysis of vehicular mobility prediction, the work should involve the following:

Simple prediction algorithm. A detailed analysis in [SKJH04a] shows the performance of a complex predictor based on data compression technique against a simple Markovian model thereby establishing an *order-2* Markovian model to perform the best. Also [Kru08] proves the effectiveness of a simple Markovian model in vehicular mobility prediction at crossroads. Hence, for a large-scale prediction analysis like the one on the Koln road network, the use of simpler learning models is desirable.

Size of past histories. A prediction model built with fewer historical knowledge will lead to a weaker relation between entities, which fail to predict accurately when current situation is more diverse to that of historical knowledge. However, more historical data lead to a relational model which exhibits balanced relation between entities, hence failing to recognize the diversity of current situation. Early works on mobility prediction have always based their prediction on huge model databases built from available mobility history. But referring to the discussion on vehicular mobility dynamics in [UF12], maintaining one large database seems unnecessary and adds to computational complexity and database update delay. This paves the way for a lightweight system that can quickly reconfigure and is available for prediction over a short time horizon.

Large-scale realistic vehicular mobility. Most of the real-world datasets employed in earlier works either describe the mobility of a small group of individuals or selected point-of-interests or commutes carried out by public transports that follow regular routes. Such limited diversity in mobility history often exposes the model to regular type of mobility patterns and fails to predict future locations for earlier unseen patterns. Also, the small geographical areas

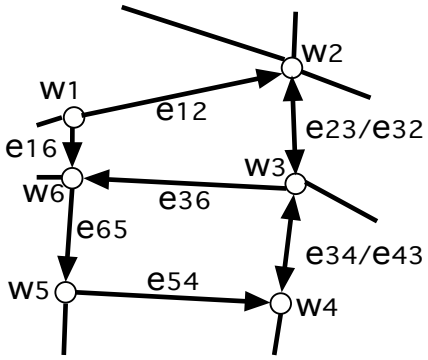

Figure 5.2: Sample network representation.

considered do not allow a spatiotemporal analysis of mobility, which is significant for network planning in wireless network or for smart traffic management.

Macroscopic prediction. Vehicular mobility privacy is also crucial along with the comfort provided by the recent technological advancements. Most vehicular mobility predictions discussed above involve learning and predicting individual mobility behaviors. Hence, there is a need for an approach that profiles users based on their similar feature and are capable of predicting future locations for these groups of users. Such an approach result in a simpler system as per the objective and preserves user privacy.

5.3 Macroscopic vehicular mobility prediction models

In this Section, we present our prediction technique that leverages the Markovian and the Sequence Matching models to quantify the predictability of vehicular user mobility in large-scale networks.

5.3.1 Markovian model

Markovian models describe a system in terms of the transition behavior among its states. When studying the mobility prediction in vehicular environments, these states correspond to the possible locations of a vehicle. To derive a discrete set of locations, mobility information is recorded when a vehicle visits a specific position on a well defined network. In the following, we consider two different network scenarios; i.e., road networks where states map to junctions and cellular radio access networks where states map to cells.

Formally, let us consider the network to be represented by a directed graph $G = (W, E)$. The graph consists of a set of vertices $W = \{w_i \mid i \in 1...N\}$, and a set of edges $E = \{e_{ij} = (w_i, w_j) \mid w_i, w_j \in W\}$ that link pairs of vertices. For the time being, let us assume

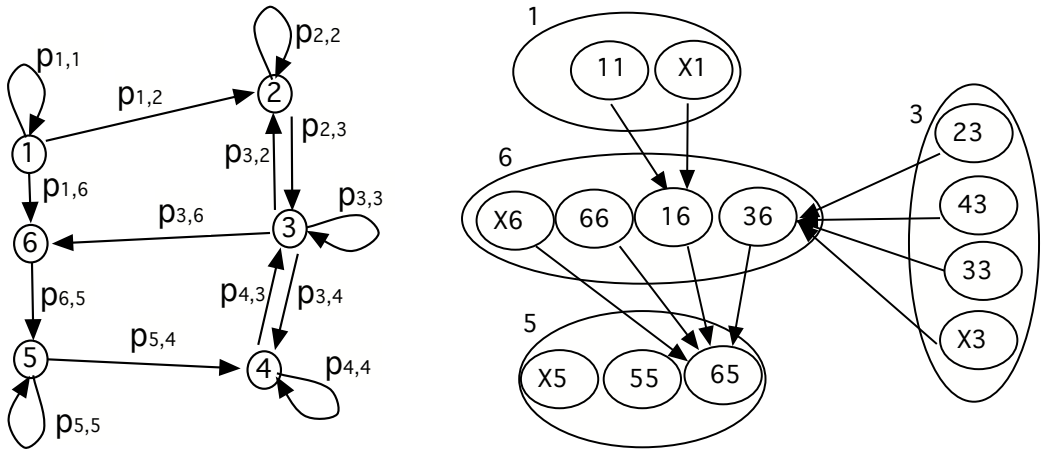

(a) Order-1 (b) Order-2

Figure 5.3: State transition diagram of order-1 and a subset of the order-2 Markovian models.

that each vertex w_i of this network also corresponds to a state in the Markovian model s_i. Then, each edge e_{ij} indicates the possibility of a transition between states s_i and s_j. An example of such a system is shown in Fig. 5.2.

The mobility of a generic vehicle across such a network can be expressed as a sequence of states $s_i^k(t)$, meaning that the vehicle k is at a state s_i at discrete time t. In order to represent the system as a Markovian model, we must assume that for every state s_i at t, the future state at $t+1$ is only dependent on the current state, i.e.,

$$P\big(s_{i+1}^k(t+1) \mid s_i^k(t), s_{i-1}^k(t-1), s_{i-2}(t-2), \cdots, s_{i-n}^k(t-n)\big) = P\big(s_{i+1}^k(t+1) \mid s_i^k(t)\big) \quad (5.1)$$

where $s_{i-1}^k(t-1), s_{i-2}^k(t-2), \cdots, s_{i-n}^k(t-n)$ are the past states or past mobility footprint. Eq. 5.1, implies that, for time-homogeneous Markovian model, the 1-step transition probability depends on state s_{i+1} and s_i but is the same at all times t; hence the terminology *time-homogeneous*. For the sake of clarity, in the following we will drop the vehicle identifier, and refer to a generic vehicle. For a time-homogeneous Markovian model, 1-step transition probability is denoted as

$$p_{i,j} = P\big(s_i \mid s_j\big),\ i,j \in 1...N \quad (5.2)$$

In such a system, there are N^2 1-step transition probabilities $p_{i,j}$. It is convenient to show them in a $N \times N$ transition probability matrix form as shown below:

$$\bar{P_1} = \begin{pmatrix} p_{1,1} & p_{1,2} & \cdots & p_{1,N} \\ p_{2,1} & p_{2,2} & \cdots & p_{2,N} \\ \vdots & \vdots & \ddots & \vdots \\ p_{N,1} & p_{N,2} & \cdots & p_{N,N} \end{pmatrix} \tag{5.3}$$

where the rows correspond to the starting state and the columns correspond to the ending state of a transition. In terms of vehicular locations, the probability of going from location 2 to location 3 in 1-step is stored in row number 2 and column number 3. We recognize such a system as order-1 Markovian model and is depicted in Fig. 5.3(a).

On the other hand, we can relax the assumption on the independency between past and the future mobility by considering an order-n Markovian model, where a future state is dependent on the current position plus a number n of previously traversed locations. A single state is then formed by encoding n previously travelled locations, hence the notation order-n. Initially we provide an order-2 model to relate our discussion to order-1 example presented above. For an order-2 model, the 1-step transition probability can be denoted as

$$p_{ij,jk} = P\big(s_{ij} \mid s_{jk}\big),\ i,j,k \in 1...N \tag{5.4}$$

Fig. 5.3(b) shows a subset of the Fig. 5.3(a) where a state s_6 in order-1 is now state s_{16}, s_{36}, s_{66}, s_{X6} in order-2. In general, an order-n transition matrix can be symbolically put together as shown in Eq.(5.5)

$$\bar{P_n} = \begin{pmatrix} p_{1...N,1...N} & p_{1...N,2...N} & \cdots & p_{1...N,N...N} \\ p_{1...N,2...N} & p_{1...N,2...N} & \cdots & p_{1...N,2...N} \\ p_{1...N,N...N} & p_{1...N,N...N} & \cdots & p_{1...N,N...N} \\ \vdots & \vdots & \ddots & \vdots \\ p_{2...N,1...N} & p_{2...N,1...N} & \cdots & p_{2...N,1...N} \\ p_{2...N,2...N} & p_{2...N,2...N} & \cdots & p_{2...N,2...N} \\ p_{2...N,N...N} & p_{2...N,N...N} & \cdots & p_{2...N,N...N} \\ \vdots & \vdots & \ddots & \vdots \\ p_{N...N,1...N} & p_{N...N,1...N} & \cdots & p_{N...N,1...N} \\ p_{N...N,2...N} & p_{N...N,2...N} & \cdots & p_{N...N,2...N} \\ p_{N...N,N...N} & p_{N...N,N...N} & \cdots & p_{N...N,N...N} \end{pmatrix} \tag{5.5}$$

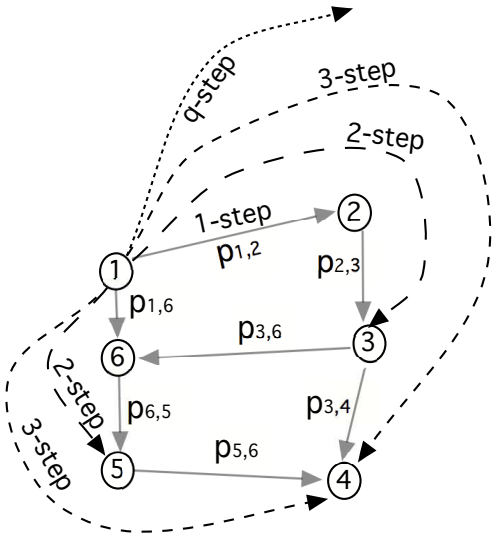

Figure 5.4: q-step prediction representation.

Higher order Markovian models with states encoded with past visited positions provide a sense of direction of mobility of a vehicle that makes it appealing for high prediction accuracy. However, as we move from order 1 to n, we narrow the usability of the model from being more general to more specific. To provide an insight into the effect of higher order Markovian chain, let us concentrate our discussion on state 6 in Fig. 5.3(a). V_6 is the vehicular volume at state s_6 at time t, $p_{6,5}$ is the transition probability from s_6 to s_5, $p_{6,X}$ is the transition probability from s_6 to any other state s_X. In order-1, volume of vehicles at s_6 is the sum of vehicles arriving from s_1 and s_3, whereas in order-2, each state is representative of one past network node in addition to the current one. Hence, Fig. 5.3(b) shows the different states, s_{16} with vehicular volume from s_1; s_{36} with vehicular volume from s_3; state s_{66} with vehicle starting at s_6; and s_{X6} with vehicles from other external connections (not shown). The growing number of states with higher order Markovian models paves the road for states with reduced vehicular transition volume per state: Such a condition often weakens the prediction outcome making it more biased. We present the effect of such condition on prediction accuracy in more detail in the following sections.

While the order of the model determines the amount of past history considered to predict the future, we are also interested in defining how far into the future the vehicular mobility is predictable with our model. The Markovian models defined above allow for 1-step mobility predictions as well as q-step predictions.

To examine the prediction with the Markovian model under the assumption that the given vehicular flow starts in a certain state s_i at time t, we simply choose $\pi(t)$ to be the probability vector with i^{th} entry equal to 1 and all other entries equal to 0.

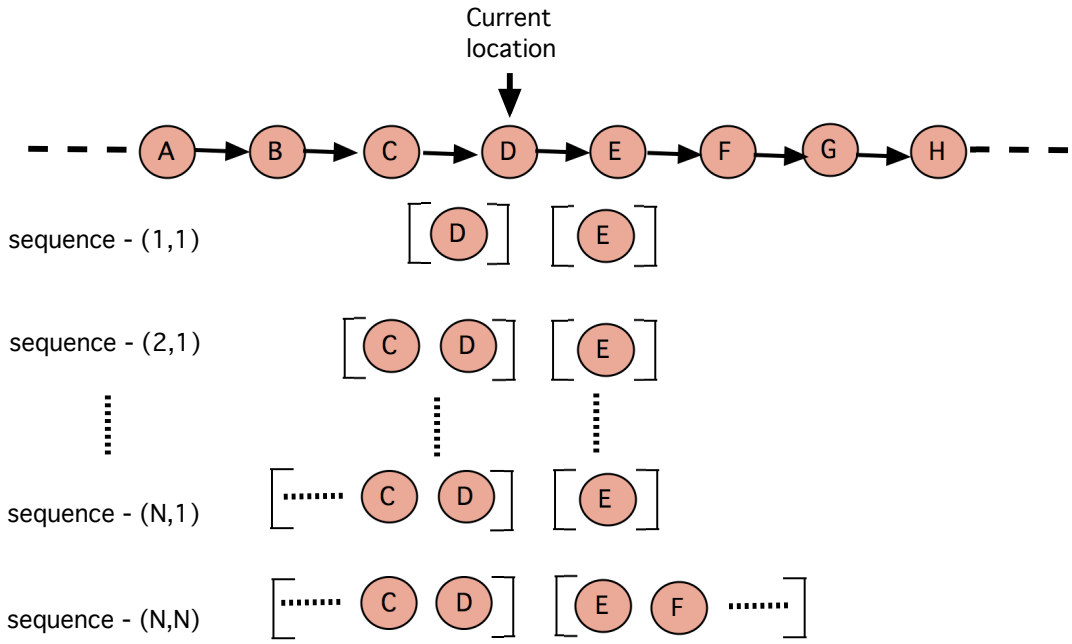

Figure 5.5: Example of sequence matching.

With the 1-step, order-n transition probability $\bar{P}_n$ defined earlier, a q-step mobility prediction for the q-step transition events that takes place at a future time $t+q$ can be obtained from Eq.(5.6)

$$\bar{\pi}(t+q) = \bar{\pi}(t)\bar{P}_n{}^q \tag{5.6}$$

Such prediction defines the extent of mobility across network nodes from a given network node or from a given location of a vehicle, hence we term q mentioned above as *prediction horizon*. A depiction of q-step prediction is shown in Fig. 5.4. In order-1, a 1-step transition from state s_1 leads to state s_2 and s_6 with probability $p_{1,2}$ and $p_{1,6}$ respectively, 2-step transition leads to s_3 and s_5 with probability $p_{1,2} \times p_{2,3}$ and $p_{1,6} \times p_{6,5}$ respectively. Similar calculations can lead us to q steps ahead.

5.3.2 Sequence Matching model

Human mobility trips, either made by public transport, walk or by personal vehicle, can be broken down into small sequences of specific locations. Modeling a system that establishes probabilistic correlation between these sequences and anticipates future sequences defines the sequence matching concept. Many variants of this simpler version of sequence matching exist, but we employ the basic version derived by modifying the Markovian model for vehicular trips.

Let $s_i^k(t)$ be the current location of a vehicle, $[s_{i-1}^k(t-1), s_{i-2}^k(t-2), ...s_{i-n}^k(t-n)]$ be past traversed footprints and $[s_{i+1}^k(t+1), s_{i+2}^k(t+2), ...s_{i+n}^k(t+n)]$ be the future locations. For the ease of explanation, we make use of the same notation defined in Markovian model and we describe the model as *sequence*-(n, q), where n is length of the past sequences and q is the length of future sequences. Also by referring to order-n in Sequence Matching, we define 1-step prediction, sequence-(n,1) unless q-step is specifically mentioned. Fig. 5.5 illustrates how a sequence matching model from a mobility footprints *"ABCDEFGH"* are constructed and used to anticipate the future sequences at D. Such mobility footprints are obtained from mobility transition on a network as shown in Fig. 5.2. To perform a prediction from the current position at D, we use the variable length of past histories n, which makes this model synonymous to 1-step, order-n Markovian model.

In 1-step prediction process, both the models have the same behaviour and the complexity. But when q is increased in q-step prediction, a line is drawn with regard to complexity and memory consumption. In Sequence Matching, an attempt is required to establish relations between already computed n length sequences and new sequences of different length $q > 1$. This process demands additional computation to derive relations between sequences and also additional memory, whereas in Markovian model, this objective is achieved by simple matrix multiplication with no additional memory requirement. Hence we extend our analysis with Markovian model and later quantify our finding against the Sequence Matching technique.

5.3.3 Online prediction with fallback

As discussed earlier, the amount of state transition knowledge used to build a relational model is crucial first step in a prediction process. One approach for training involves maintaining and updating one large database. Though such model may be more knowledgeable and anticipate future states given current or past states, it is often impractical in a real world road network featuring large-scale mobility transition data, due to processing and access delays.

Motivated by the results from [UF12], observing that vehicular mobility pattern changes frequently with time, we define an *online* approach for training. We construct a training database with a time interval window $[t - T_B, t]$ and this information is used to anticipate the vehicular mobility in the period $\{[t, t + T_F] \mid T_B, T_F \in T\}$. The prediction is always carried out with the most recent training interval hence the challenge lies in deciding an ideal interval for training and prediction. Fig. 5.6 shows a sliding window which moves across the mobility data constructing the model and predicting the future vehicular mobility for the

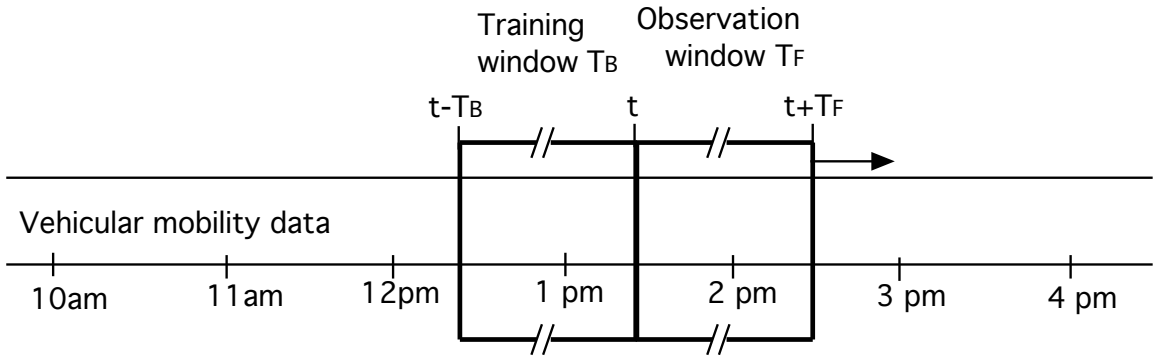

Figure 5.6: Sliding window of training and observation.

interval $[t, t+T_F]$. We also adopt a cyclic approach, where a training database created during the last time interval in the complete mobility dataset is used to predict the beginning of the mobility data. The impact of window sizes T_B and T_F is discussed in Sec. 5.4.2.

Prediction during order-n Markovian model or the sequence-(n,q) in Sequence Matching, with $n > 1$ often suffers from *prediction failure*, a situation where the model is not knowledgeable about the current state because it was never encountered during the training phase. This situation is addressed by a technique called *fallback* [Las06, DK04]. According to this approach, a transition probability matrix is maintained for every Markovian order and Sequence Matching length n from 1 till n_{max}, given an observed state during the prediction period, the fallback algorithm starts by checking for the corresponding state (of length n) in training databases. In case of failure, it falls back and checks in the lower order databases until either it finds a future prediction or it has tried a first order prediction and has still failed. This technique helps to evaluate and select the optimal state length to build a model that encounters less failures as well as delivers best prediction accuracy.

5.3.4 Estimating the prediction accuracy

Following the online prediction, monitoring the performance of the model also plays a significant role in the prediction process. As previously stated, we concentrate on predicting the macroscopic vehicular flows i.e., our prediction is an estimation in terms the portion of vehicular flows that transit from one state to the next. Hence a similarity metric is used to calculate the error between the prediction based on training and the reality. We employ the *Euclidean distance* for this purpose for its low computational complexity.

In general, the prediction error Δ at a state i at time interval $[t + T_F]$ is obtained as

$$\Delta_i^{[t+T_F]} = \sqrt{\sum_{j=1}^{M_i} \left(p_{i,j}^{[t,t+T_B]} - p_{i,j}^{[t,t+T_F]} \right)^2} \tag{5.7}$$

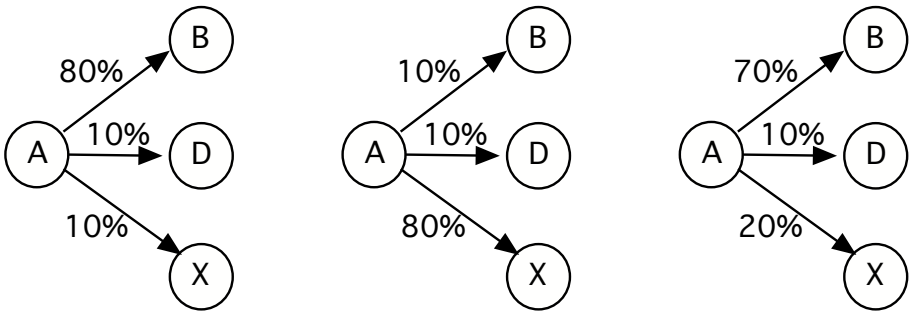

(a) Training (b) Observation 1 (c) Observation 2

Figure 5.7: Simple order-1 showing state transition with different traffic behaviors.

where $p_{i,j}^{[t,t+T_B]}$ is the transition probability as predicted using the training knowledge to that of the observed reality $p_{i,j}^{[t,t+T_F]}$, M_i is the total transiting states form a given state i. If $v_i^{[t,t+T_F]}$ is the vehicular volume i.e., the number of vehicles transiting by state i during $[t, t+T_F]$ then overall weighted error of state i over all time durations over the day is given by

$$\Delta_i = \frac{\sum \Delta_i^{[t,t+T_F]} \times v_i^{[t,t+T_F]}}{v_i}, \quad where\ v_i = \sum v_i^{[t,t+T_F]} \tag{5.8}$$

Similarly, overall error for the whole system is obtained by

$$\Delta = \frac{\sum_{i=1}^{N} \Delta_i \times v_i}{\sum_{i=1}^{N} v_i} \tag{5.9}$$

Fig. 5.7 shows a state transition diagram with order-1 states. To demonstrate the scenarios of high and low prediction error, let us consider the traffic dispersion from state A as shown in Fig. 5.7(a) and we use this information to build our relational model. Predicting the traffic at some future time interval with this relational model yields a high error of 98.99% using the similarity metric in Eq. (5.7), if the traffic flow in the reality follows the pattern shown in Fig. 5.7(b). Such condition is often seen when the relation model is used to predict a traffic flow that is contrary to which the relational model is built upon. Also this situation is more likely to occur in case where the relational model is defined considering low number of transitions between states, which mostly is the condition during the early period of the morning. On the other hand, as seen in Fig. 5.7(c) the traffic pattern is similar to that of the

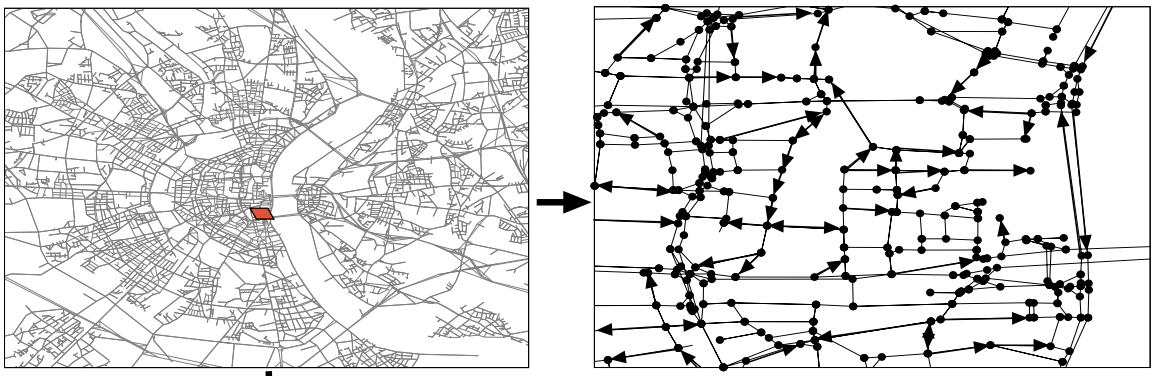

(a) Crossroads as network node

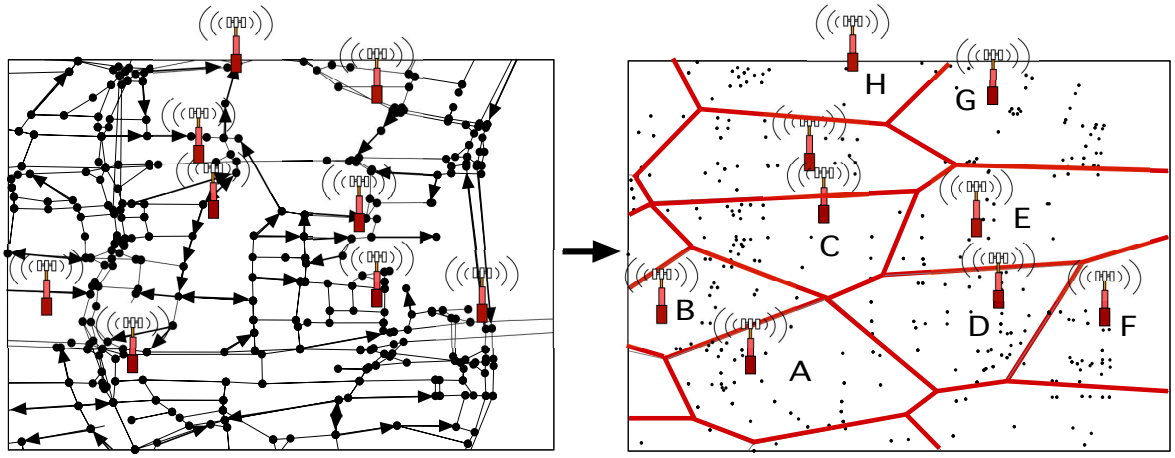

(b) Cellular cells as network node

Figure 5.8: Road network and cellular network representation.

training, this yields a high prediction accuracy of 85.85%. This condition is more anticipated during the peak hours of the day as the relational model is built upon greater number of traffic definitions.

5.4 Scenarios

In this section, we discuss how the prediction models discussed above apply to two types of networks i.e., road network and cellular RAN. To that end, we first need to adapt our Koln vehicular mobility description to suit the models. Later, we perform the calibration to decide the model parameters that provide the best prediction accuracy.

5.4.1 Vehicular mobility adaptation for prediction model.

To employ the Koln synthetic vehicular mobility dataset in our prediction model, we adapt it to as described below. In real-world datasets, vehicular mobility information could be derived from sources like on-board GPS module [JHP+03, DPPW10, LWY+06], road side detectors [iTe11, WK99, PDP11], induction loops [BK09] or even by processing cellular data records [LM96]. However in such data it is laborious to obtain high granularity and also inconsistencies in the quality of the data requires heavy post processing and approximations. Since we use synthetic dataset, we have the advantage that we can easily obtain high position granularity i.e., every second. In case of a prediction model based on the road network, we reduce the granularity of our synthetic dataset to road intersection level by capturing the moment when vehicle change the road segment along their trip and a complete vehicular trip can be represented as

$$\bar{W}^k = \{w_i^k(t) \mid \forall i \in N\} \tag{5.10}$$

On the other hand, in the case where the prediction model is applied to the cellular network, we leverage the Voronoi tessellation discussed in 4.2.1, and adapt the vehicular mobility across the RAN infrastructure. The vehicular mobility is now represented as a series of base station identifiers and is defined as

$$\bar{C}^k = \{c_l^k(t) \mid \forall l \in C\} \tag{5.11}$$

These base stations cover a specific geographic area and hence covers a portion of the road network or in other words, each base station covers a small subset of the road intersections, also $c_l \subseteq W$. In Fig. 5.8(a), we consider a small area of the city center (in red) to show the road network constituted of crossroads represented as dots and the road segments. In Fig 5.8(b), we place the base stations on the road network and each polygon in red mimics a cellular coverage covering number of road intersections (black dots), hence the vehicular volume in any cell is thus an aggregation of volume across these crossroads. In cellular network these base stations are interconnected to provide uninterrupted service, hence vehicular users at any given time are connected to the core network through at least one base station. Studying the level of predictability of vehicular users with our approach in these two networks helps to understand the prominence of various parameters in mobility prediction.

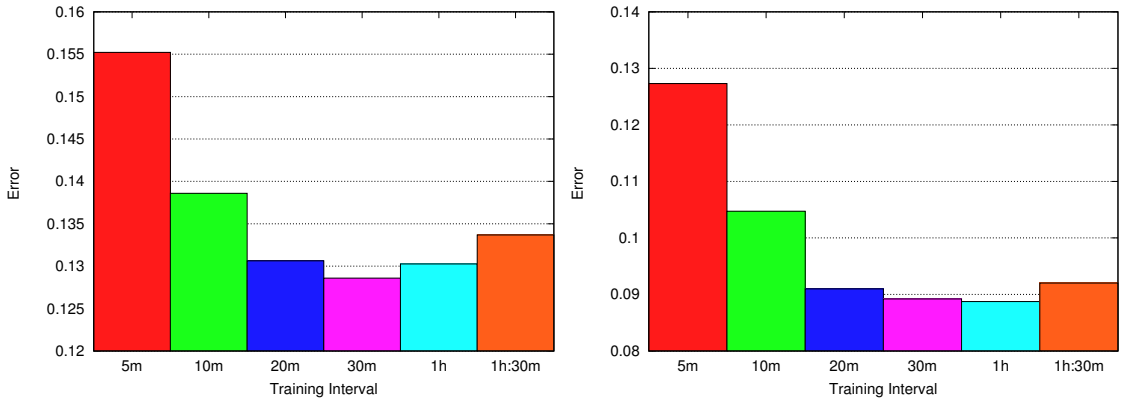

(a) Fixed observation of 5 minutes

(b) Fixed bservation of 30 minutes

Figure 5.9: Road network: Varying training interval T_B

5.4.2 Impact of model parameters for crossroads

In this section, we explore the model parameters affecting the macroscopic online prediction of vehicular mobility so as to decide which provide the best prediction accuracy. We remark that all the calibration is performed on the Markovian model, and the results are later verified with the Sequence Matching model in Sec. 5.5.

Training interval T_B

Driven by the objective to build an online prediction model with a fixed length window of mobility information, deciding on the optimal duration of the training and observation time intervals becomes crucial.

Initially, we run through different training intervals with the observation interval T_F fixed to 5 minutes. This observation interval exposes the models sensitivity towards the vehicular volume at crossroads, hence a training interval that performs best with 5 minutes observation interval is likely to perform better with greater T_F. Similarly we start our analysis with order-1 model, a state in this order exhibits no favorable direction information unlike higher order states that can aid the prediction process.

Fig. 5.9(a) portrays the effect of varying the value of the training interval T_B. The prediction error is calculated as discussed in Sec. 5.3.4. High error is seen with lower values of T_B, this behavior is expected as these periods restrict the states of the model to capture fewer mobility transitions that create bias among the transiting states, hence leading to greater dissimilarities. But as we increase T_B, the error improves, indicating large number of states

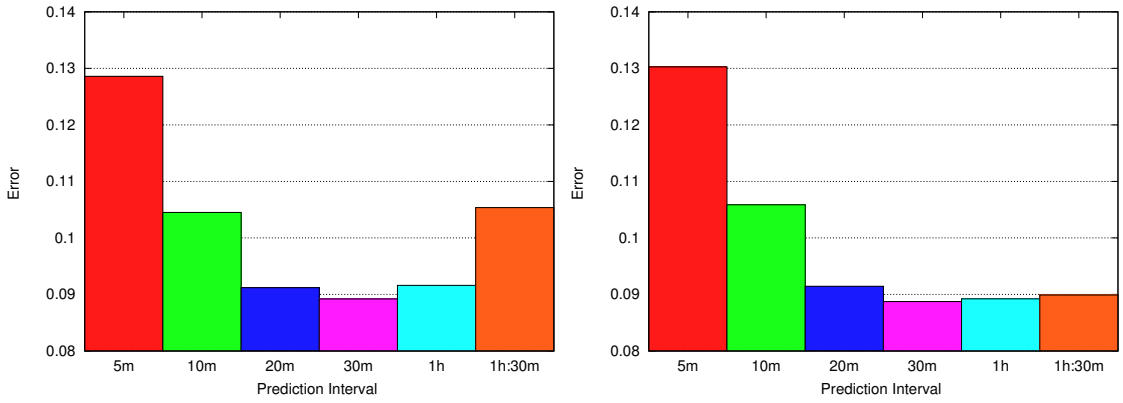

(a) Fixed training 30 minutes

(b) Fixed training 60 minutes

Figure 5.10: Road network: Varying observation intervals.

with better volume and thereby fairly presenting the prominence among the transiting states. A surprising outcome in our analysis is that the prediction error tends to grow beyond 30 minutes training interval, which indicates that aggregating over long time intervals to acquire more knowledge narrows in fact the ability to distinguish different mobility patterns thereby failing to recognize the significant transition behaviors.

Though these prediction errors with higher training intervals seem to be visually greater, the actual percentage varies from 9% to 15%. Before concluding with the lowest error, we repeat our analysis with a larger T_F of 30 minutes to pick the optimal candidates for our future analysis. Comparing the results in Fig. 5.9 leave us with training interval T_B equal to 30 and 60 minute as the best options where most of the mobility transition are expected to be captured.

Observation interval T_F

The observation interval defines how often the training database is updated, and shorter intervals often induce processing delays, whereas larger intervals lead to wider gaps between mobility patterns in training and observation. In an attempt to find the best observation interval, we continue our analysis with the best candidates of T_B i.e., once with 30 minute training interval and later with 60 minute training interval.

Fig. 5.10 reports similar trend as seen in Fig. 5.9. Although with the best training interval of 30 and 60 minutes, higher error is seen with observation interval of 5 minutes. This re-assures us that capturing mobility transitions at a state that span over shorter time interval serves no good for any reliable prediction. Since the observation interval immediately follows the training interval T_B, we speculate that the change in mobility patterns in these intervals

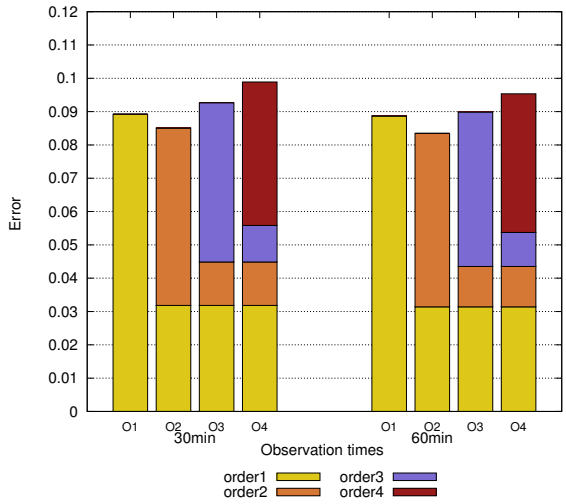

Figure 5.11: Road network: Overall error of 60 minutes training with different Markovian order.

is unlikely to be the reason for high error. As we will later see, investigating the relation between the volume at a state with error provides a deeper understanding. As we have to decide on the T_F configuration, we select 60 minute as our final training interval, 30 and 60 minutes as the observation intervals for our future analysis.

Prediction model order

The order-1 model employed in the previous sections establishes a benchmark for the prediction errors. But this could be revised by considering the higher order models, since a state in higher order model is often encoded with the sense of direction of mobility within each state. These higher order states narrows the number of possible future transiting states compared with order-1 where the options are omnidirectional. Also, as we increase the order of the model, we encode additional past mobility information into each state thereby increasing the confidence of prediction, but at the same time we make the model more specific to smaller road traffic flows' mobility patterns, which increases the possibility of prediction failure. The fallback technique helps us to compare the results obtained with every order by maintaining equal number of predictable cases in all orders. In order to quantify this, we proceed with 60 minute training interval and 30 minute observation interval as concluded in the previous section.

Fig. 5.11, depicts the impact of the order of the Markovian model on the prediction error. Each bar with On is used to imply the corresponding order n of the Markovian chain. As we adopt a fallback technique to overcome the prediction failure condition, the error contributed

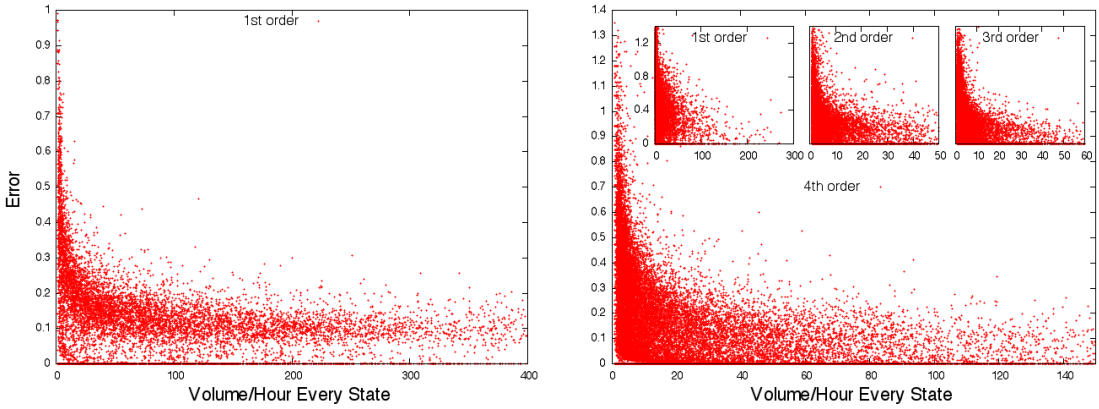

(a) Order 1

(b) Order 4

Figure 5.12: Road network: Error vs volume-per-state/h over 24 hours calculated for every state.

by the respective lower orders is denoted with different colors within each bar. We see that order-2 performs better than order-1 as best accuracy is expected from higher orders, whereas error grows with order-3 and order-4 which is contrary to our expectation, hence we stop our analysis at order 4.

To investigate this behavior, we present the scatter plot in Fig. 5.12, where each dot represents a state. The x-axis defines the average vehicular volume transiting through a state (hence providing the volume-per-state) over 24 hours and the y-axis defines the prediction error. We present only the order-1 and order-4 plots as order-2 and order-3 follow the same trend. The inset figure in Fig. 5.12(b) shows the states in the lower order which contribute to prediction error incase of prediction failure. A correlation is seen with decreasing error as the volume-per-state increases, and this behavior is followed in all the orders and in their fallback lower orders. This confirms our earlier speculation that a lower volume-per-state is one of the factors that lead to higher error as a result of inability to distinguish the prominence between the transiting states.

We also perform the analysis with the 60 minute observation interval as shown using block of bars in the right of Fig. 5.11, and we see scarce improvement over 30 minute T_F. Given the results in this section, we will consider 60 minute T_B and T_F for our spatiotemporal analysis.

5.4.3 Impact of model parameters for the cellular network scenario

We now consider the vehicular mobility transition across cells of the RAN, and present how the prediction model parameters differ from those presented in the study of road network case.

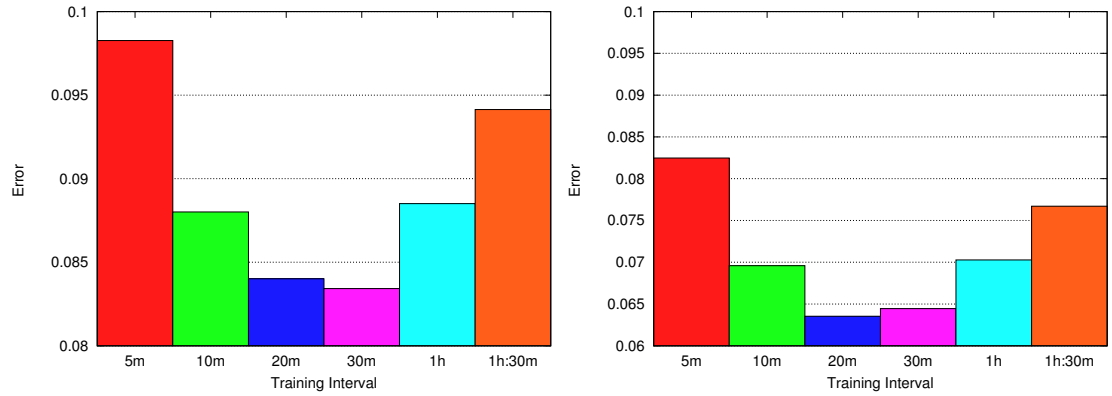

(a) Fixed observation of 5 minutes

(b) Fixed observation of 30 minutes

Figure 5.13: Cellular network: Varying training intervals.

Training interval T_B

Following the same procedure as discussed in Sec. 5.4.2, with observation interval of 5 minutes and order-1 Markovian model, we study the effect of varying training interval of mobility information on prediction error.

Fig. 5.13, show different training intervals T_B along the x-axis and the y-axis defines the prediction error. Comparing the results to the equivalent plots in case of road network, we see that the worst case error peaks to 9.7 to 9.8%, which falls close to the lowest prediction error in road network scenario. This result is expected as in RAN, the volume-per-state is considerably higher than that of a state in road network, as the volume across number of crossroads belong to a single cell in RAN and the transitions between RAN states are thus greater in number. This in-turn allows the model to distinguish significant transitions between states at lower training interval leading to lower error compared to the road network scenario. But this advantage do not last long as the model saturates failing with a greater error after training interval of 30 minutes. We confirm this result with a slightly wider observation interval of 30 minutes and thus choose 20 minutes and 30 minutes training intervals for our future analysis.

Observation interval T_F

To determine the best observation interval T_F for cellular network, we perform our analysis with 20 minutes and 30 minutes training intervals with order-1 Markovian model. Fig. 5.14, shows the prediction error in the y-axis with varying observation intervals along the x-axis. In

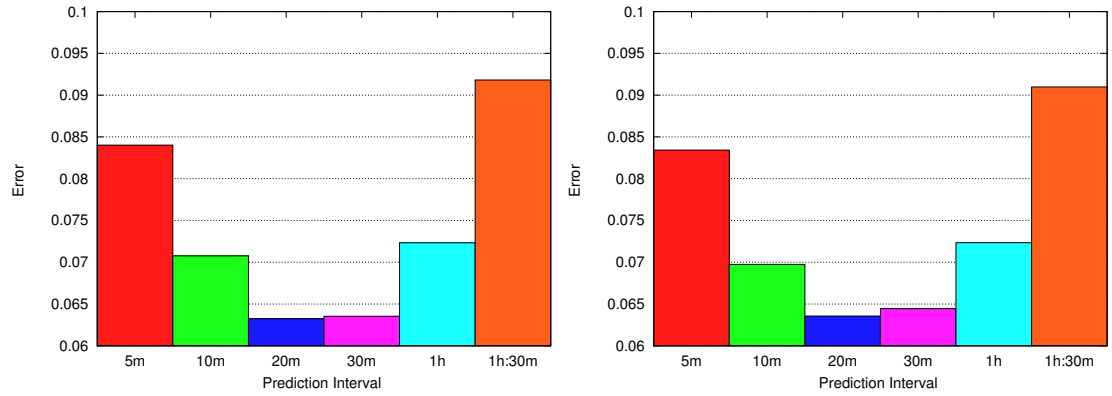

(a) Fixed training 20 minutes

(b) Fixed training 30 minutes

Figure 5.14: Cellular network: Varying prediction intervals.

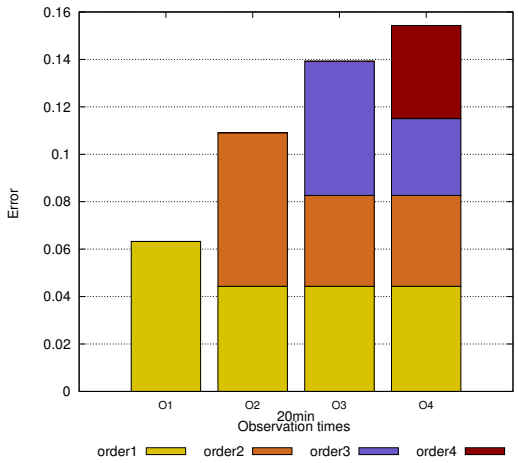

Figure 5.15: Cellular network: Overall error with 20 minutes training and observation intervals.

both cases of T_B, 20 minutes and 30 minutes of T_F performs the best, this forms the optimal window sizes with which the model can better anticipate the future macroscopic traffic flows. But as we increase the window size, the model seem to be less effective as more and more mobility information with diverse traffic flows increases the dissimilarity between the training and observation patterns.

Prediction model order

In a road network, a crossroad can have a maximum of four to five directly connected intersections, in an ideal hexagonal representation of the cellular network a cell can have six

connecting cells, whereas in a Voronoi tessellation, a cell can have a varying number of neighbors which exceeds the maximum number of connections in former cases. This can have a direct reflection on the prediction error when the order of the model is increased.

Fig. 5.15, shows the prediction error at each order of the model represented as a bar, with the error contributed by the lower order fallback states shown in different colors. Since there is a minimal difference between the 20 minute and 30 minute observation intervals we present the best results i.e., 20 minutes. We see that the order-1 model performs the best and is better than the best prediction obtained with the order-2 model in road network scenario. But the error increases more rapidly with higher orders and performs worst than that seen in road network scenario. This happens because of the greater number of neighboring cells that give rise to a greater combination of states, thereby leading to large number of lower volume states. Also we see that the error contributed by order-4 states is much lower to that compared with lower order fallback states this also hints that during the 20 minutes of T_F and T_B the vehicles can travel to fewer cells which leads to lower number of higher order states.

5.5 Performance evaluation

After the calibration of the parameters affecting the macroscopic flow prediction in Sec. 5.4, we use these parameters to study the spatiotemporal dynamics and the maximum horizon of the prediction models.

5.5.1 Road network scenario

Prediction horizon

Referring back to the discussion in Sec. 5.3, we present the results of the q-step prediction with the Markovian and the Sequence Matching techniques.

In Fig. 5.16, the x-axis shows the different orders of the Markovian chain and each colour indicates a future step q. The sub-plot in the first row highlights the percentage of dissimilarities between the training and observation, i.e., the prediction error, while the sub-plot in the second row shows the amount of vehicular volume that was accounted for during the 1-step prediction but are unavailable for subsequent q-step predictions, because the corresponding vehicles finish their journey within the q-steps. The third row explains the average volume-per-state for q-steps and for every order of the prediction model. The sub-plot in the fourth row gives the average number of future possible locations or paths that can result from the prediction at a given state. The 1-step prediction results are the same as shown in Fig. 5.11.

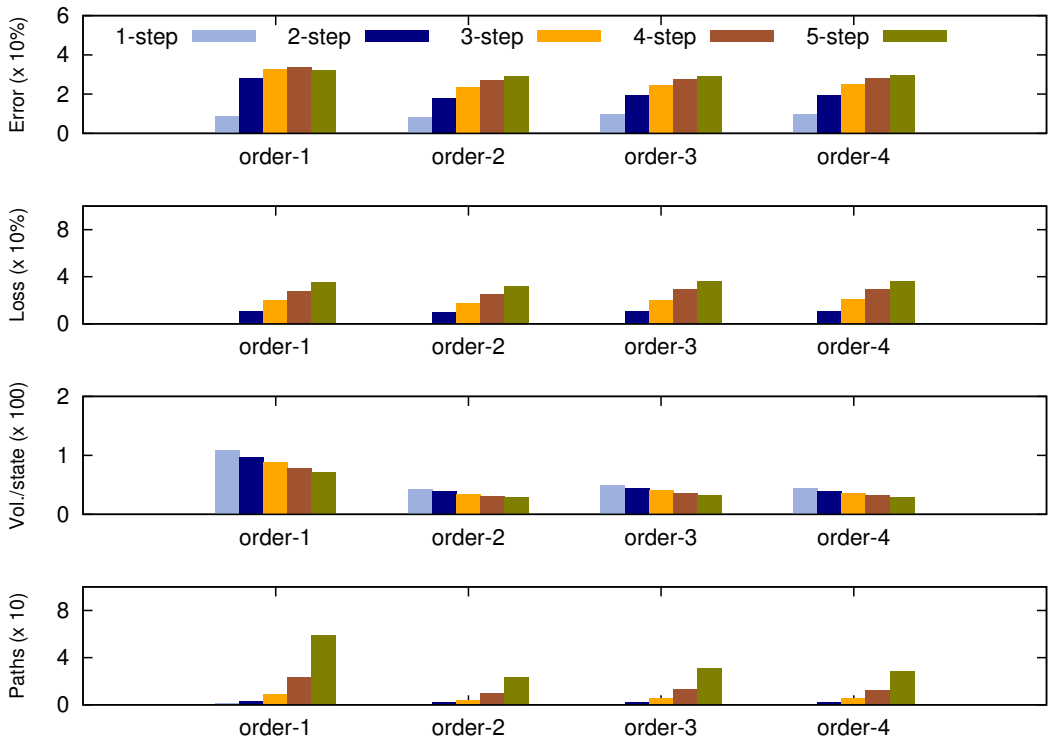

Figure 5.16: Road network: q-step mobility prediction using Markovian model.

We see that as we go beyond 1-step prediction, the prediction error increases. This is because of the fact that as we predict farther into future, the number of directions or paths a user can take increases, this is evident from the sub-plot in the fourth row. Due to the large number of possible next location, the prediction becomes less accurate or less reliable. In case of 1-step prediction, the maximum number of future locations is limited to the number of outgoing connections from the current crossroad, hence is more predictable than when $q > 1$. Now, as we increase the order of the Markovian state and compare the outcome of the respective q-step predictions, the order-2 model tends to perform well comparatively with the order-1 results, but no significant improvement is seen in order-3 and -4. Since the error for q-steps in the first sub-plot is calculated over different volumes, it becomes necessary to consider the amount of volume which are unavailable in the subsequent q-steps which we define as the loss percentage. If this percentage is considered as error, then we can envision that the prediction error at order-1 constitutes to higher error compared with other orders. Overall, we can say that the vehicular volume is predictable to 1-step with 85% to 90% accuracy, whereas the prediction model is not much reliable to predict 5 crossroads ahead as it provides an accuracy around 30%.

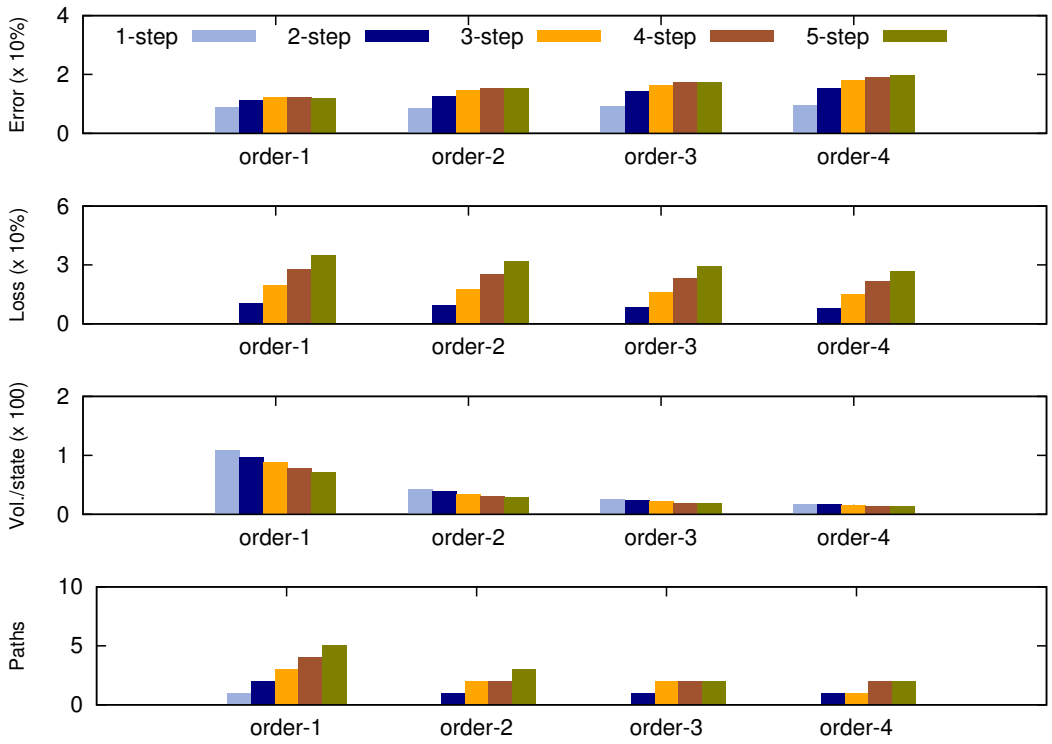

Figure 5.17: Road network: q-step mobility prediction using Sequence Matching.

Earlier we advocated that having higher volume-per-state often yields a good prediction, but comparing the first row and third row the results seem to contradict, with higher orders along x-axis having much lower volume-per-state and also reduced error compared with order-1. The reasons we can think is that the states in order-1 have no knowledge about the past mobility information, hence the options of future locations are much greater in number as can be seen in fourth row. Also in order-1, the transition probability from order-1 state to a future q-step state is more susceptible to alteration by the vehicular flow probabilities from the intermediate states, whereas this does not affect the higher order state because of the encoded mobility history in each state.

Fig. 5.17, shows the q-step behavior for the Sequence Matching model with varying sequence length and the results in each row hold the same meaning as discussed for the Markovian model. We see that as we increase the step q, an increased error is observed, which clearly associates the reason to lower volume-per-state and more paths to be predicted as seen in third and fourth row respectively. Also the prediction error increases with the order comparing the respective steps, due to the increased number of states with order and to the lower volume-per-state. On the other hand, in the Markovian model, the prediction error improved with order. But if we look at the percentage error, the prediction using Sequence Matching yield

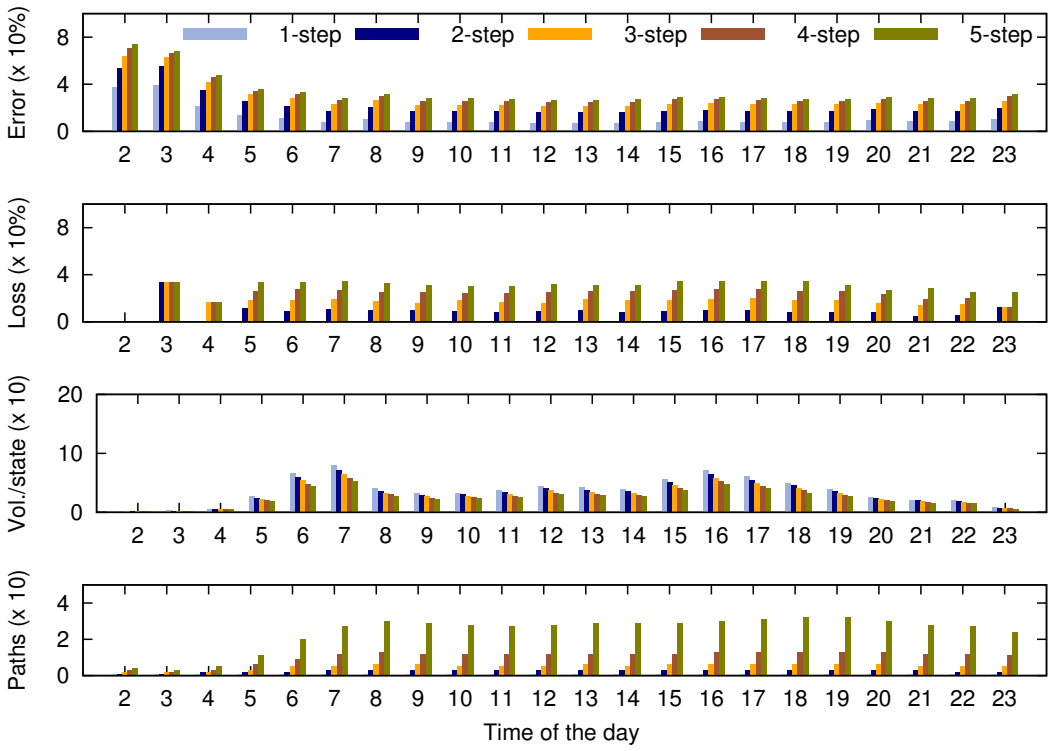

Figure 5.18: Road network: Variation of prediction error with time of the day with order-2 Markovian prediction.

s around 15% in all steps in order-1 which is far more better than 35% to 38% in Markovian model. Considering the loss percentage with prediction error, we can say that around 50% of vehicular volume is predictable up to 5 junctions ahead which is more impressive than the Markovian model. But as previously said this improvement comes at the processing and memory requirements needed to build the q-step model with Sequence Matching technique. In the Markovian model a simple matrix multiplication is used to compute the future step transitions, which gives rise to more possible paths is not followed in reality, hence contributing to the error and making the prediction less reliable.

Spatiotemporal analysis

Fig. 5.18 portrays the 1-step prediction results obtained every hour with rows defining the prediction error, loss percentage of vehicular volume, volume-per-state and number of future locations from a state similar to that defined earlier in Fig. 5.16. The x-axis shows the time of the day; overall these plots illustrate how these values vary with time. We concentrate our discussion on the configuration that provide best prediction as concluded in the previous sections. To recall from our discussion in Chapter 3, the vehicular traffic volume in our Koln dataset is the highest at around 7 am and 5 pm, i.e., the rush hours of a typical day. We

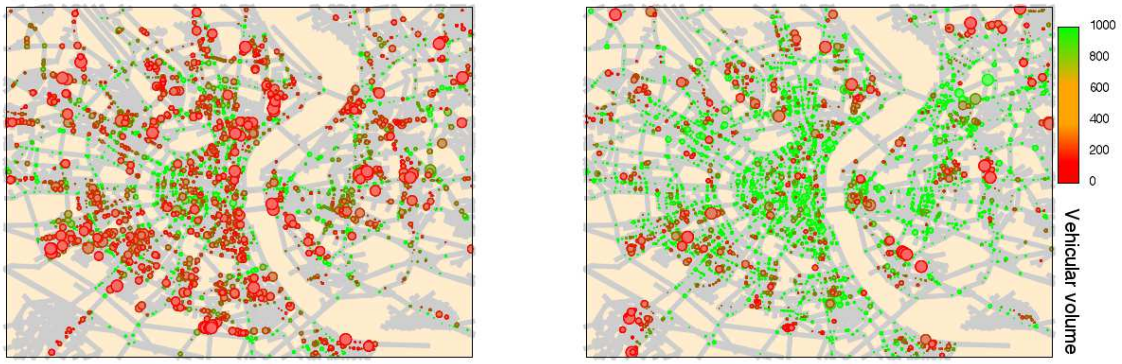

(a) 5 am

(b) 6 am

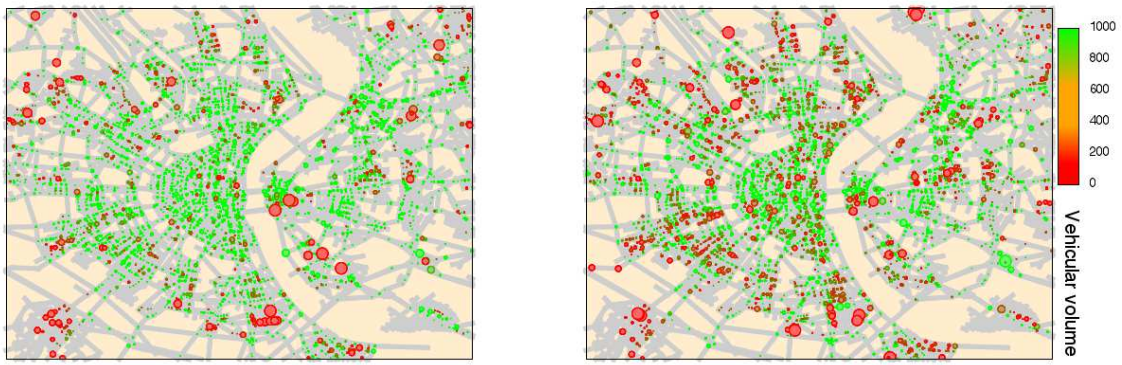

(c) 7 am

(d) 8 am

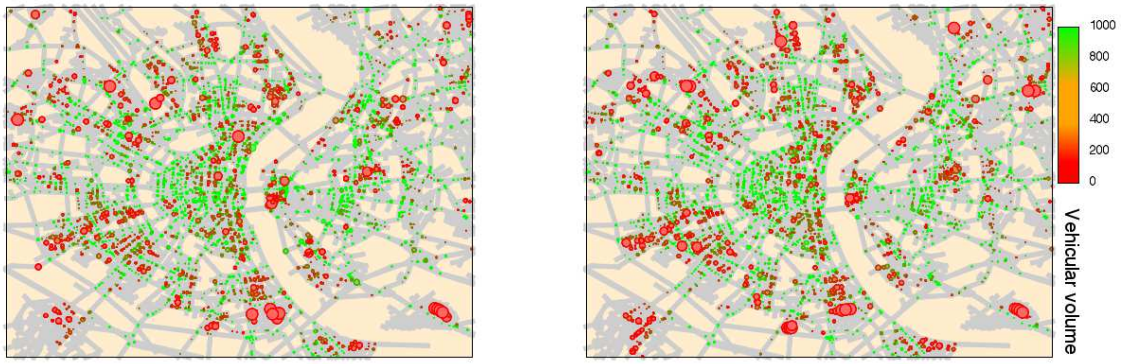

(e) 9 am

(f) 10 am

Figure 5.19: Road network: Spatiotemporal distribution of order-2 prediction error and vehicular volume (5 am to 10 am).

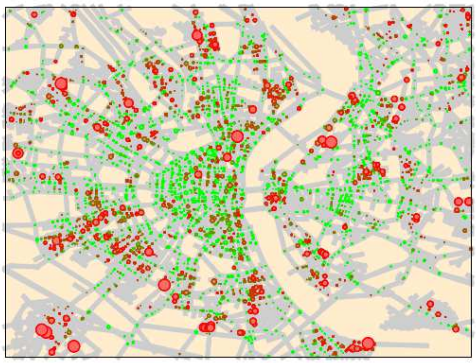

(g) 11 am

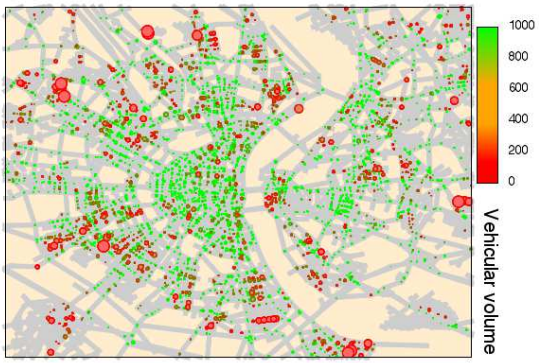

(h) 12 pm

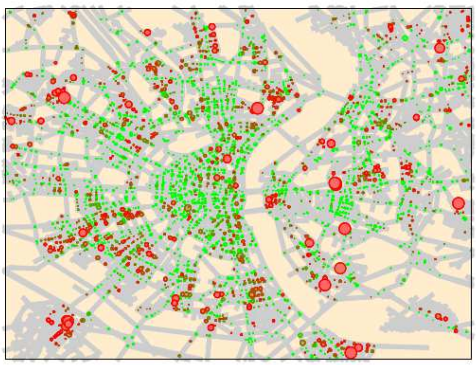

(i) 1 pm

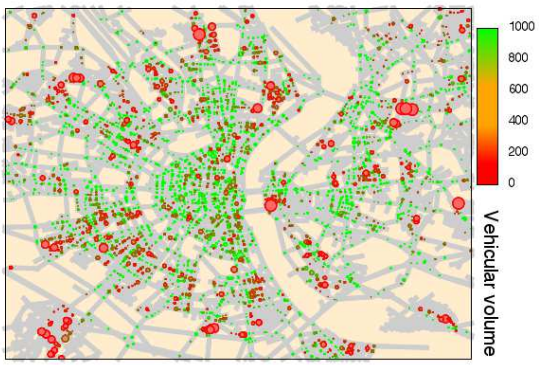

(j) 2 pm

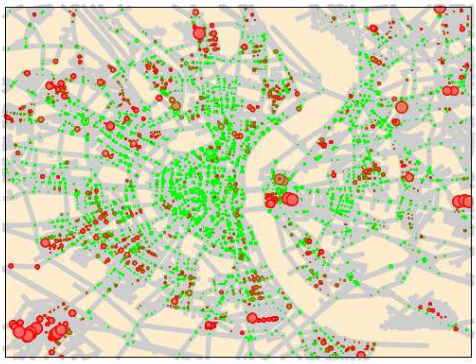

(k) 3 pm

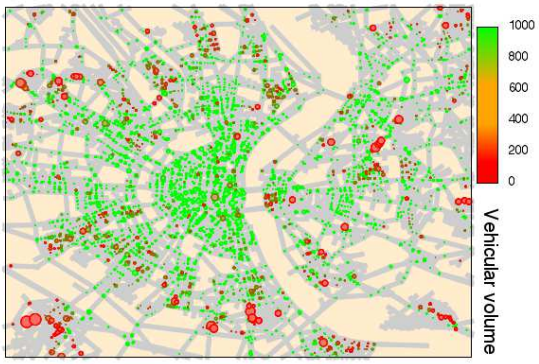

(l) 4 pm

Figure 5.19: Road network: Spatiotemporal distribution of order-2 prediction error and vehicular volume (11 am to 4 pm)

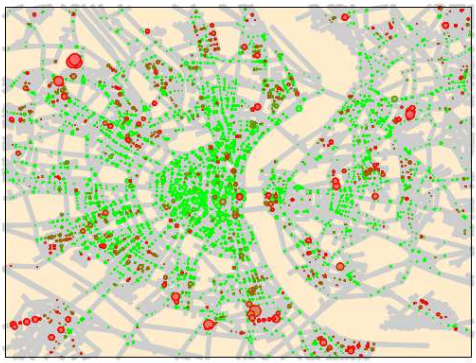

(m) 5 pm

(n) 6 pm

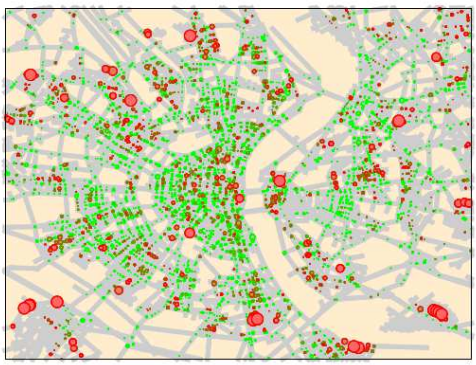

(o) 7 pm

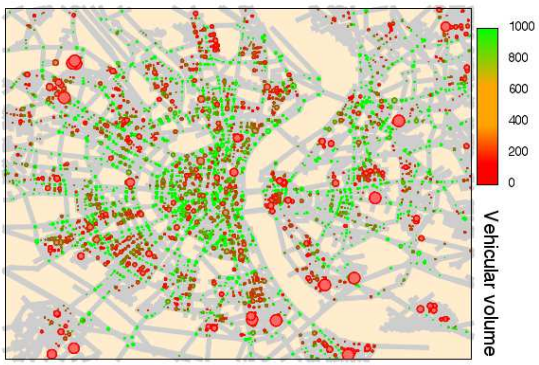

(p) 8 pm

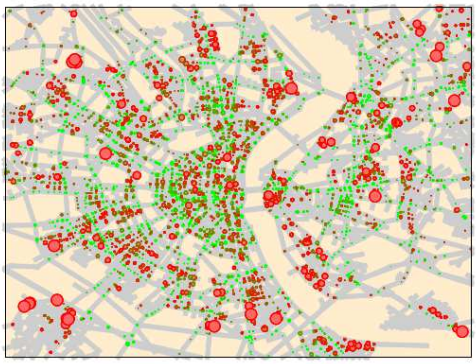

(q) 9 pm

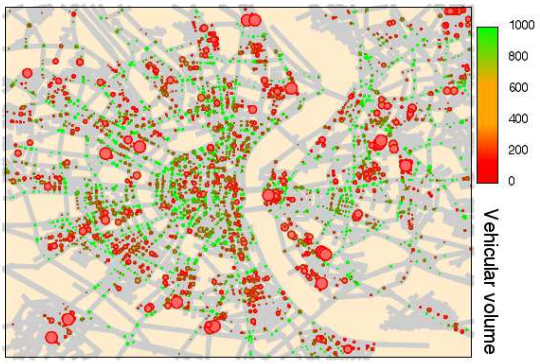

(r) 10 pm

Figure 5.19: Road network: Spatiotemporal distribution of order-2 prediction error and vehicular volume (5 pm to 10 pm)

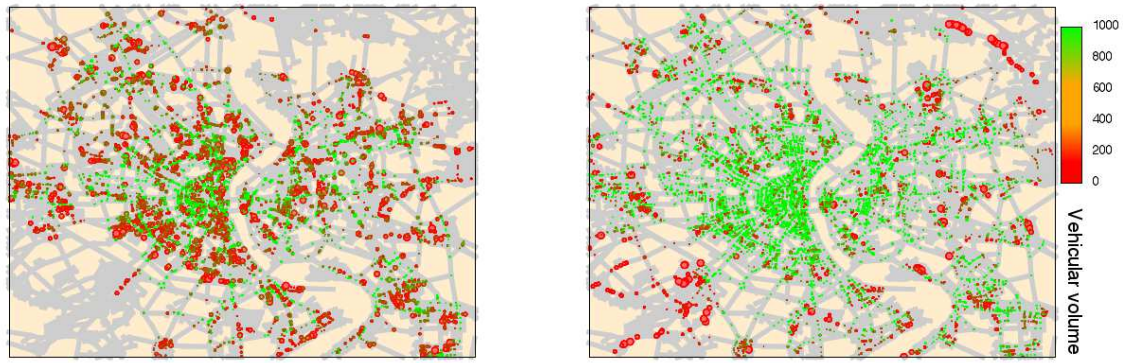

(a) 5 am 1-step

(b) 6 pm 1-step

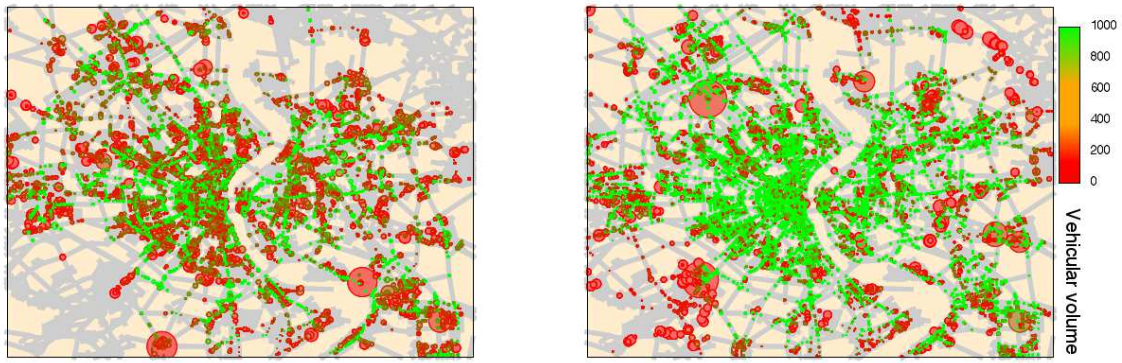

(c) 5 am 5-step

(d) 6 pm 5-step

Figure 5.20: Road network: Comparison of 1-step and 5-step prediction results using Markovian model.

do not see any prediction error during the early 1 hour period, as the mobility dataset lacks the mobility information between midnight and 1 am hence model cannot be trained for prediction. Also we cannot predict during 1 am and 2 am as this data is used for training to predict vehicular flows between 2 am and 3 am. Moreover this early morning hours account to a small number of vehicular flows, thus the model built on such data exhibits a higher error when compared with rest of the cases during the day. The model performs the best when predicting from 7 am to 8 am: one could think two possible reason for this behavior, either the randomness in the mobility patterns is low or the model is more knowledgeable to anticipate future mobility patterns. Fig. 4.5 explains this behavior, where morning peak hour is dominated by the vehicular flows towards the city centre, which are more predictable due to less diverse traffic behaviors than the afternoon, where mixed mobility patterns are observed.

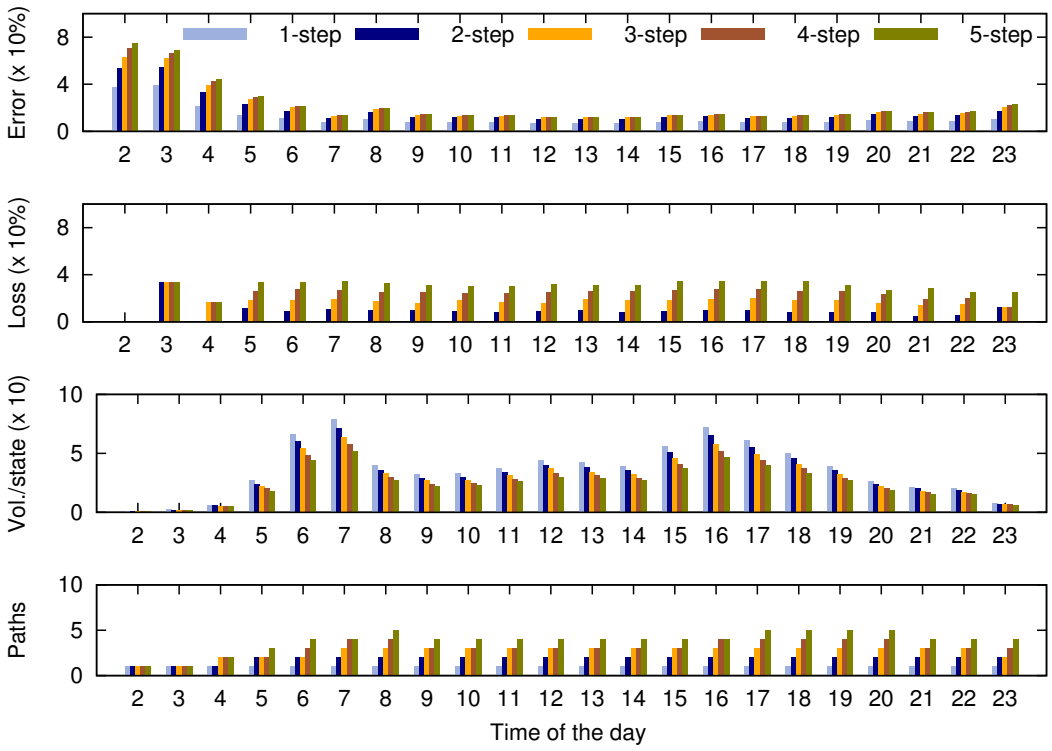

Figure 5.21: Road network: Variation of prediction error with time of the day with order-2 Sequence Matching prediction

Fig. 5.19, shows the spatial analysis of order-2 Markovian model, where each dot represents the prediction error at a crossroad, the size of the dot depicts the prediction error and the color of the dot turns from red to green when the traffic volume increases. We see that as we go from 5 am to 7 am the dots turn green, indicating that traffic volume per state grows as well as the size of the dot decreases showing reducing prediction error, this is seen in the third row of Fig. 5.18. But we can see bigger red dots at the outskirts, where the traffic generally enters the city, due to traffic congestion drivers explore different routes, thereby adding to the randomness via new patterns that may not be followed often. Comparing the peak hours of the day, 3 pm to 5 pm looks greener with good traffic volume-per-state. But concentrating at the city centre, 7 am have smaller dots than the evening hours, because the mixed mobility behavior in the evening leads to slightly higher error. Also few areas in the outskirts are always red and large dots can be seen, these are the states with lower traffic volume resulting in biased situation.

Fig. 5.18 also shows that, as we try to predict q-steps into future, the combined error i.e, when loss percentage is added to prediction error; looms between 50% to 90% in all time of the day. Higher number of states, lower traffic volume per state, and higher number of predictable future locations are the main reasons which makes the prediction unreliable.

It is quite surprising to see such a behavior at q-steps during traffic rush hours of the day, so to have a deeper understanding we compare the prediction results of 1-step with 5-step. We choose an off-peak hour in the morning, 5 am and evening peak hour, 6 pm. From Fig. 5.20, we learn that during the morning 5 am 1-step case, the error is contributed by the crossroads in the city centre which may be due to low volume per state. At 6 pm in the 1-step case, the prediction is good around city center but we can see larger red dots in the outskirts which seems to be the main source of error due to lower traffic volumes. But in the 5-step case, at both hours the error is mostly contributed by the larger red dots at the outskirts of the city, green dots in the city center are comparatively larger than those in 1-step prediction. This means that crossroads with lower traffic volumes are not the eligible candidates for higher step prediction and needs careful selection of areas where larger mobility history could be beneficial.

Fig. 5.21, shows the variability of the prediction error over the time of the day using the Sequence Matching model. As previously mentioned, the Sequence Matching model is the same as the Markovian model when both are used for 1-step prediction, and the real distinction is seen in q-step prediction. Observing the error variability over different times of the day, the trend of the error is similar to that seen in Fig. 5.18 with states at peak times predict more accurately than at off-peak hours due to more knowledge learned from diverse traffic flows. One interesting observation is that except the early period of the day, the error difference between 1-step and 5-step prediction varies around 10% when compared to around 30% in case of Markovian model. This lower value is because of the use of historic data to train the q-step model, whereas in Markovian model matrix multiplication is used. This introduces a greater number of paths or possible future locations from which an accurate prediction is hard or to achieve and moreover these paths may not be followed in the real world. Since the error difference between q-step prediction in both models is small (between 0% and 10%), it is hard to differentiate between the spatial plots, so we avoid presenting the same.

Altogether, these analysis confirms that the mobility patterns of the vehicular users can be predicted at crossroads with an acceptable accuracy of around 90% during the peak hours of the day by observing the traffic dynamics for a short period of time. In the light of these results, our approach best fits systems that are designed with low computational capability, low memory and that which intelligently operates by conserving power. This approach thus yields the same advantage as the traditional ones, but in addition can offer prediction in situations where historical data over days may not be available.

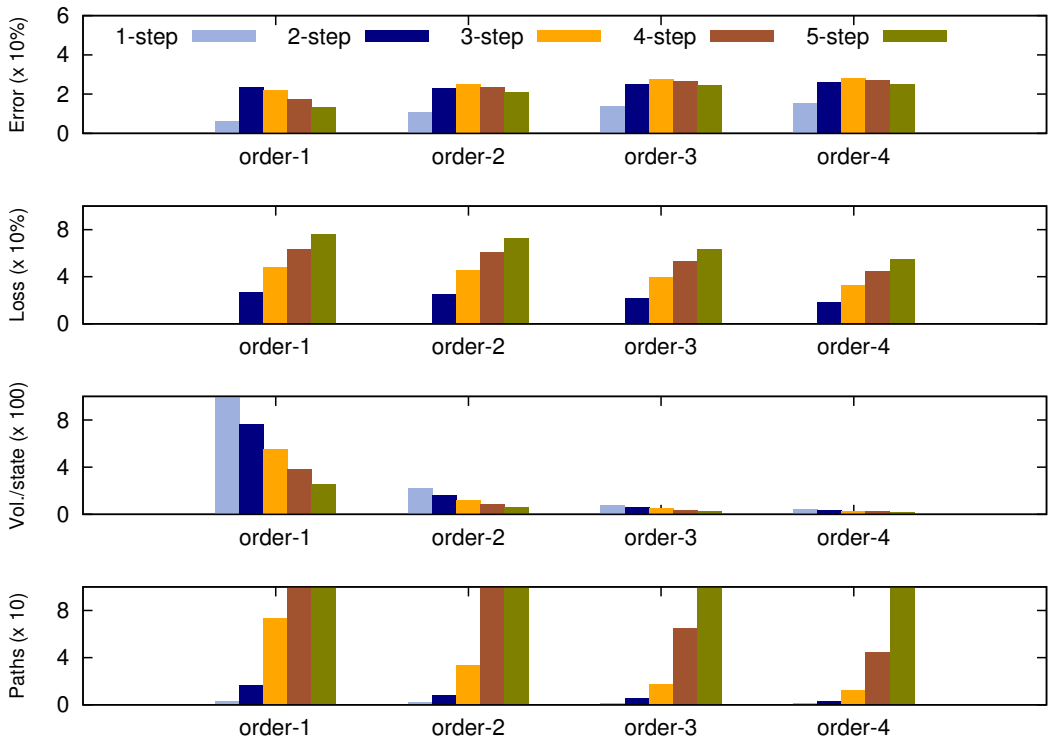

Figure 5.22: Cellular network: q-step mobility prediction using Markovian Model.

5.5.2 Cellular network

Prediction horizon

Fig. 5.22 depicts the q-step prediction in the cellular network scenario using the Markovian model. The description of the sub-plots are the same as mentioned in Sec. 5.5.1. From our earlier discussion, in cellular networks, transition probabilities are defined between base station cells and each cell covers a subset of the road network, hence the volume at a state is expected to be much larger than those in a road network scenario. This is evident from third row of the sub-plot where the volume per state is hundreds of vehicles/hour per state higher when compared to the equivalent sub-plot for the road network case. Looking at 1-step results, the error increases with order-n, hence providing the same results as discussed in Sec. 5.4.3, and this is also true when the respective steps of different orders are compared. But when we compare the results of different steps in the same order, we see that error tend to decrease, the reason behind this is that as we go farther to predict vehicular mobility up to five cells ahead, loss percentage increases. This is acceptable as with 20 minutes training and observation intervals, vehicles are capable of traversing only few cells. To compare these results, considering the loss percentage along with prediction error leaves us with a prediction accuracy of around 10% in 5-step prediction. As seen in case of road network scenario the

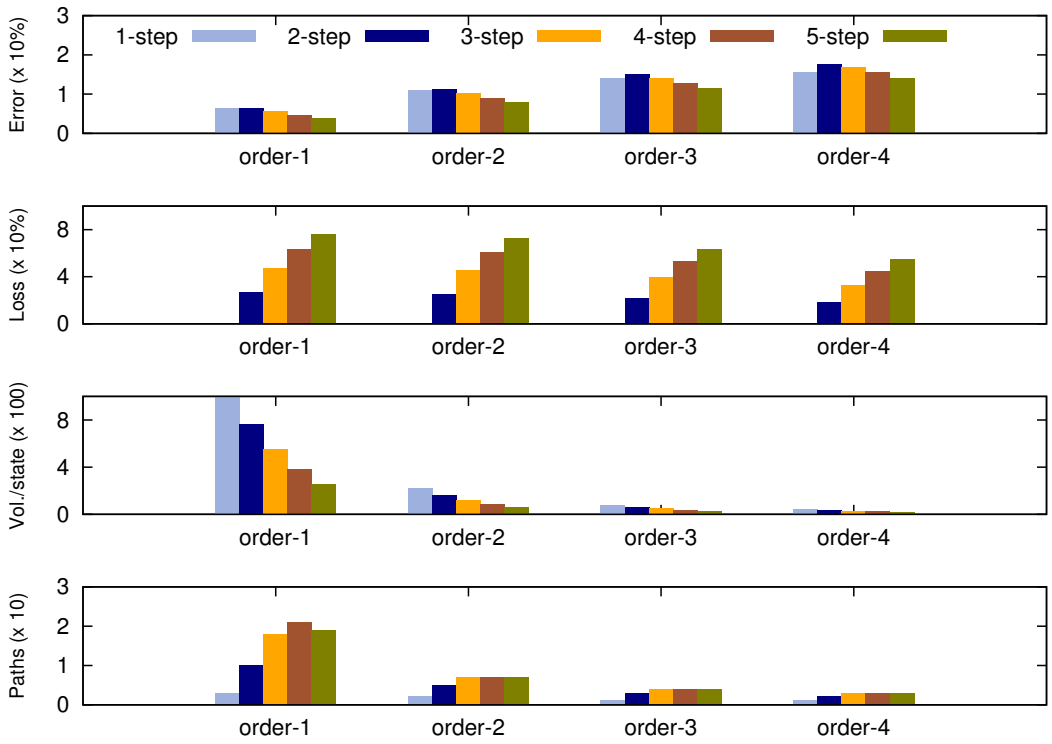

Figure 5.23: Cellular network: q-step mobility prediction using Sequence Matching.

number of possible future locations or paths are higher in order-1 compared to that of higher orders, as sense of direction information in higher order states exposes only fewer possible option. With this, one would expect to have a good prediction in higher order, but it is not the case because of the nature of the cellular coverage, where the number of neighbors of a cell is greater than the maximum number of neighbors of crossroads in road network scenario. With large number of states, and given the nature of q-step transition probability calculation in Markovian model, this lead us to unexpected types of mobility transitions which are not actually followed in reality due to the layout and constraints imposed on the mobility in road network.

Fig. 5.23 shows the q-step Sequence Matching results in cellular network coverage. The error follows similar trend as in case of road network scenario. One interesting thing to note is the number of paths as shown in the sub-plot in fourth row when compared with the Markovian model. This highlights the uncertainty of using Markovian model, whereas in Sequence Matching the q-step model is built with historic data in the interval $t - T_B$ as discussed before, hence includes only valid paths. The prediction error grows with sequence length n because of the higher number of states, but comparatively the error is less than as with Markovian model. Higher prediction accuracy of more than 90% is obtained in 1-step

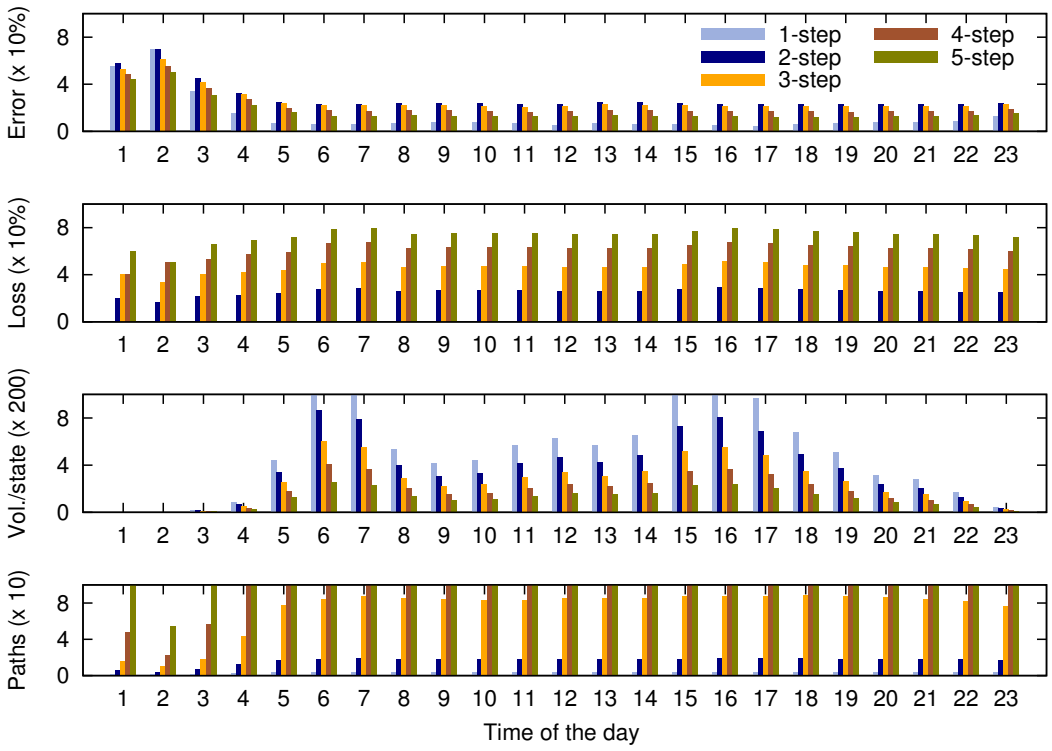

Figure 5.24: Cellular network : Variation of prediction error with time of the day with order-1 Markovian prediction.

prediction but the prediction accuracy at higher steps is not reliable with just 20% accuracy in order-1, which is best achieved. For a q-step prediction, one could also think of using a large window size $t - T_B$ so as to improve the volume-per-state and reduced q-step loss percentage and thereby the error.

Spatiotemporal analysis

As seen in Fig. 5.22, the overall error in order-1 model performs the best, hence we present an in-depth analysis of the order-1, q-step prediction error variability over the time of the day in Fig. 5.24. This illustration has the similar definition as the equivalent plots in case of road network with the x-axis presenting the time of the day and the y-axis defining different measure of value associated with prediction at a cellular cell. Now we have the prediction results at 1 am, this is because the prediction is carried out with training $t-T_B$ and observation $t + T_F$ intervals of 20 minutes hence the values presented every hour is the weighted average within that hour.

The error variability is similar to that seen in case of crossroads but the percentage loss is much greater with around 10% of vehicular flows traversing five cells ahead. If this loss is

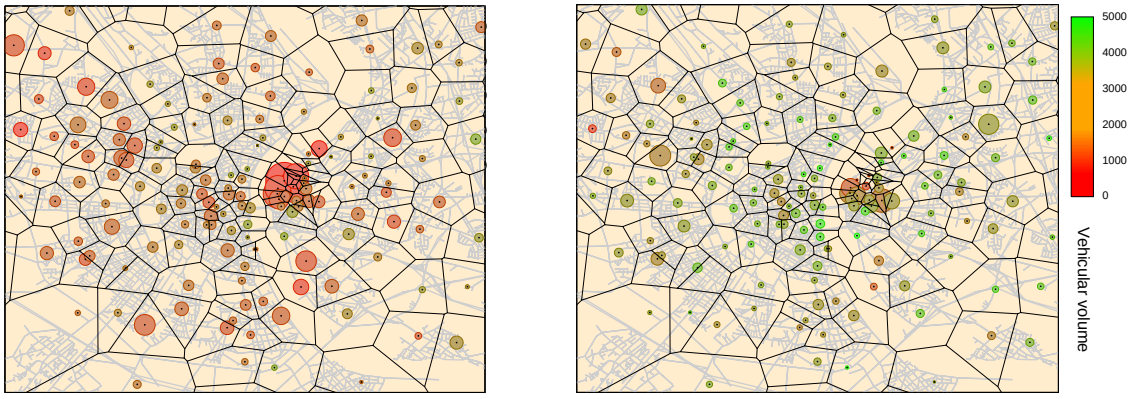

(a) 5 am

(b) 6 am

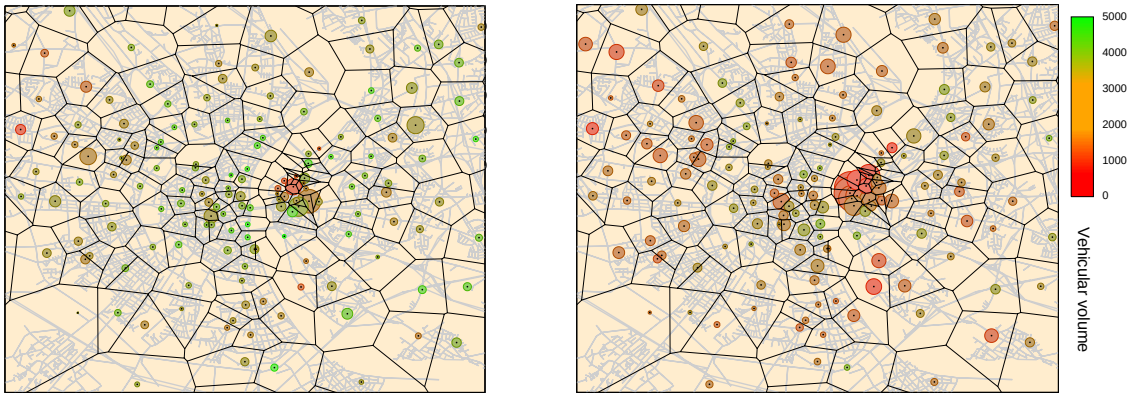

(c) 7 am

(d) 8 am

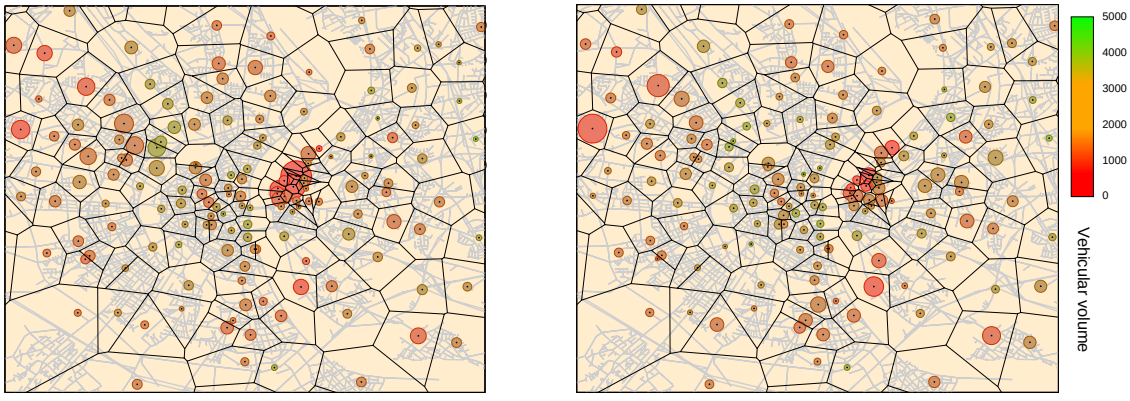

(e) 9 am

(f) 10 am

Figure 5.25: Cellular network: Spatiotemporal distribution of prediction error and vehicular volume (5 am to 10 am).

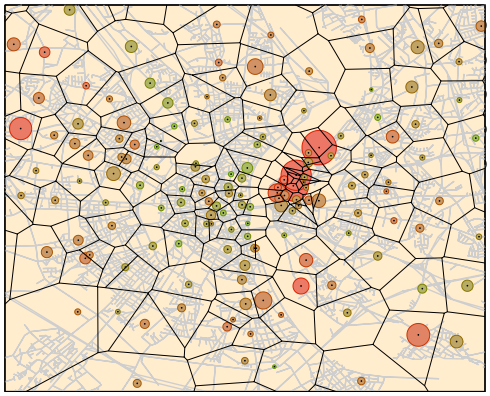

(g) 11 am

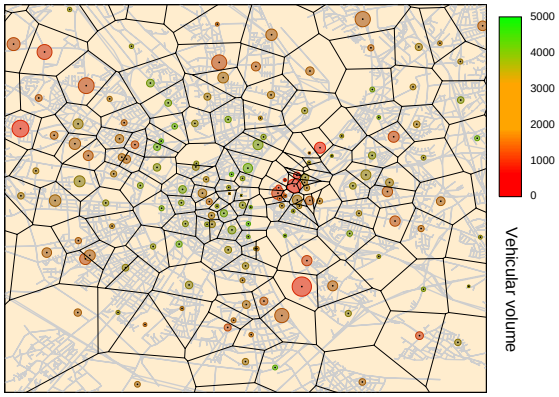

(h) 12 pm

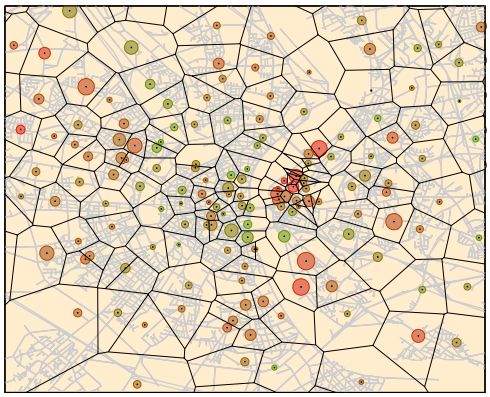

(i) 1 pm

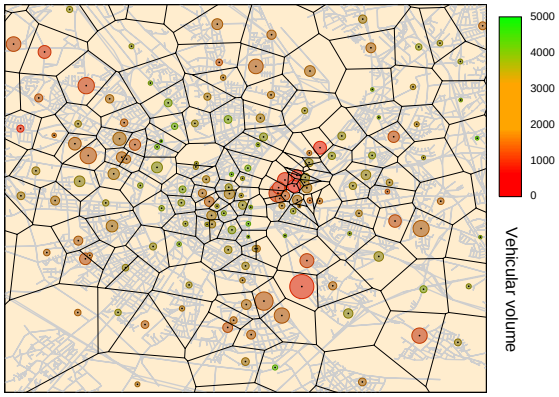

(j) 2 pm

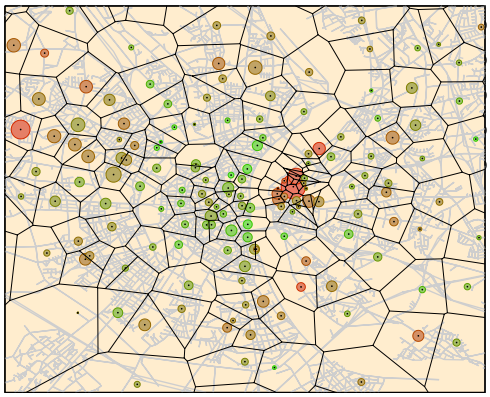

(k) 3 pm

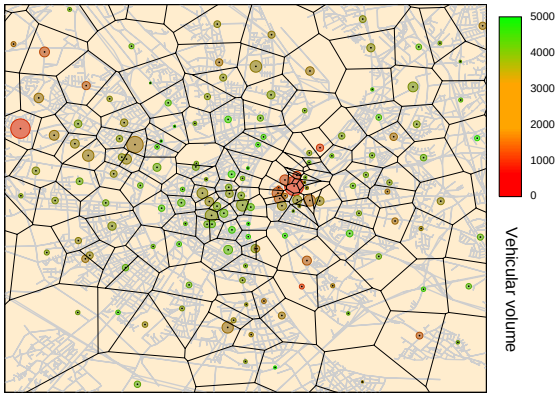

(l) 4 pm

Figure 5.25: Cellular network: Spatiotemporal distribution of prediction error and vehicular volume (11 am to 4 pm)

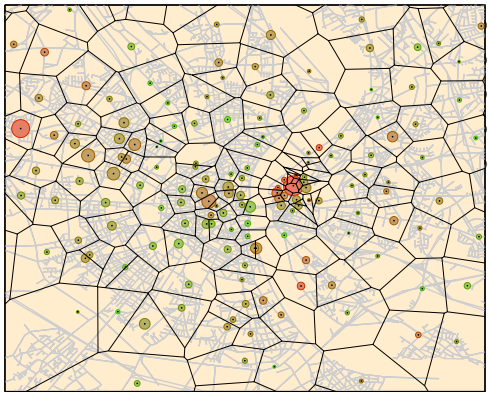

(m) 5 pm

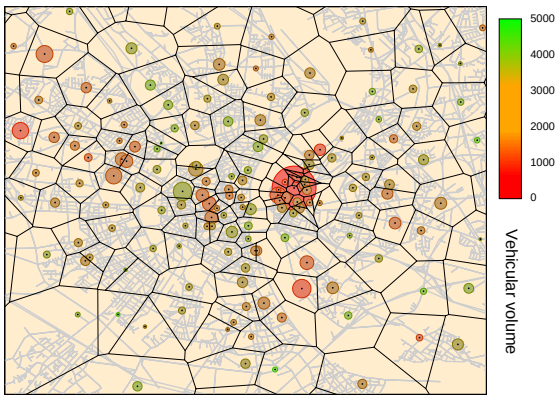

(n) 6 pm

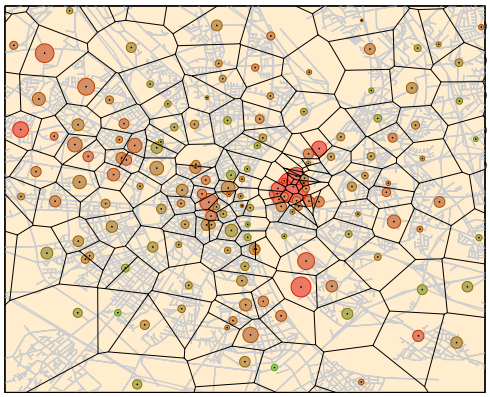

(o) 7 pm

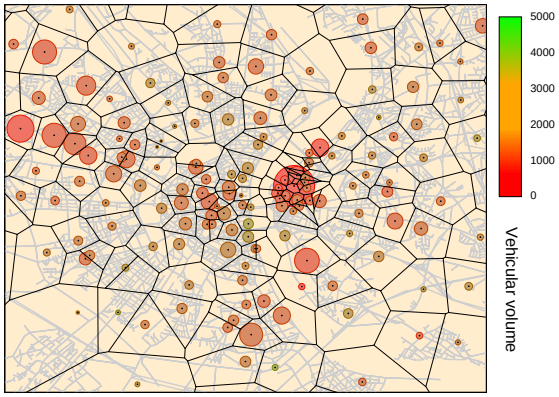

(p) 8 pm

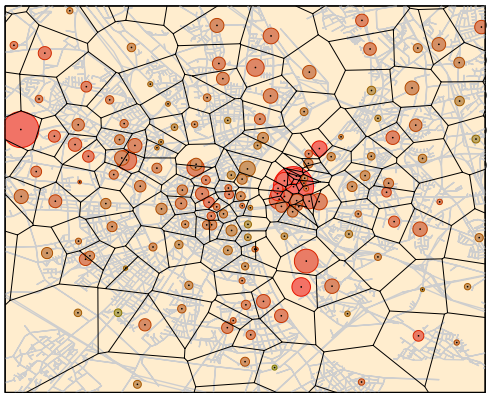

(q) 9 pm

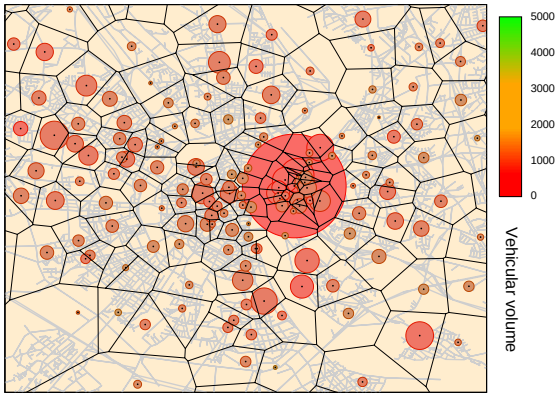

(r) 10 pm

Figure 5.25: Cellular network: Spatiotemporal distribution of prediction error and vehicular volume (5 pm to 10 pm)

considered as error then the prediction error almost reaches 100% when anticipating 5-step into future in most of the time during the day. The variability in traffic volumes is expected to be high during the peak hours and is shown in the third row and is 20 times greater when compared to the road network scenario. There is no significant change seen in the number of possible future paths but is comparably higher than the road network scenario in all time of the day.

Fig. 5.25 shows the spatial plot with Voronoi tessellation representing the cellular coverage over the geographic area. During the early hours, when the volume per state is very low the error tend to be high but during the peak time, only few cells have high prediction error in the city center shown as bigger red dots indicating higher error due to very low volume. But, if we look into the outskirts of the city, where cells cover the highways, the prediction error is always at its minimum, as the mobility pattern does not change much due to the nature of underlying road network. If we look around the city center the prediction error varies a lot with time, as a cell covers road network with multiple crossroads extending the road segments to number of cells that facilitates the dynamic mobility across cells. In Fig. 5.24, we see that the 1-step prediction error is almost equal to that of the peak hour of the day in most cases, but the spatial plot show higher error with big dots at different time snapshots, this is true but the dots signifies a very low traffic volume by its color. When we calculate the weighted average over all cells every hour, we see a minimal difference ranging less than 5% between different 1-step error values at different times. Hence spatial analysis serves the purpose of identifying these areas of the city where careful attention is required to address the effects of poor prediction accuracy.

With the percentage losses added to the q-step prediction error and lower volume per state makes it less interesting to examine the spatial analysis so we skip the same.

Fig. 5.26 presents the prediction results using the sequence matching technique. The definition of the x and y coordinates are the same as discussed for the Markovian model. It is evident that the number of future possible paths directly reflect the level of prediction accuracy, though the losses are the same as seen in the Markovian model, there is a sense of improved accuracy of around 10% to that obtained with the Markovian model.

5.6 Summary

In this Chapter, we discussed the significance of mobility prediction at crossroads as well as in the cellular radio access network architecture. The challenges associated, and the need for

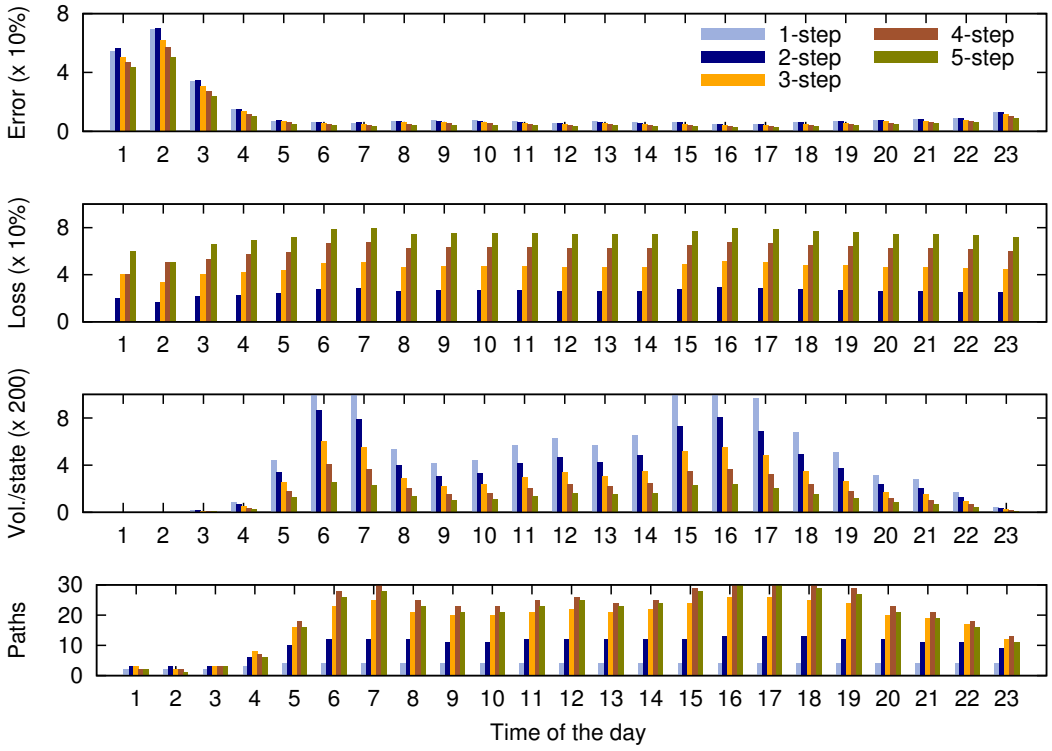

Figure 5.26: Cellular network : Variation of prediction error with time of the day with order-1 Sequence Matching prediction.

a lightweight prediction approach are discussed. The system involves Markovian chains and sequence matching techniques to predict the near future locations of the vehicular users with the knowledge of near past mobility locations traversed. The effect of vehicular mobility on the prediction model parameters is presented along with the spatiotemporal analysis of the prediction accuracy. In the road network scenario, the vehicular mobility at the city center is more predictable compared to the outskirts, hence more consideration or vehicular traffic knowledge is required to perform prediction in such cases. In cellular network scenario, cells in the outskirts have good prediction, but in the city center the prediction varies. In both cases our results let us speculate that by considering near term past mobility information the 1-step vehicular mobility patterns are predictable with 90% accuracy.

6 Conclusions and perspectives

In this Chapter, we present the conclusions drawn from the results of this thesis and we highlight the possible future extensions that could be interesting to further validate our findings.

6.1 Conclusion

In this thesis, we stress the need for understanding the vehicular dynamics in an urban environment following the incorporation of smart devices in vehicles that hint to become future communication hubs. To that end, research efforts have been concentrated on modeling and simulation of synthetic mobility datasets considering the cost and complexity involved in real-world measurements. The growing data traffic demand is luring the cellular operators to adopt small cells to accommodate the data traffic demand. Such attempts have often raised more challenges in quality of service provided, hence predicting the mobility of cellular subscribers is emerging as a need to better anticipate the mobile demand. In spite of all these developments next generation cellular service needs new approaches based on different categories of subscribers recognized by their unique mobility and access behaviors.

Through this dissertation, the following objectives are achieved:

- **Identification of the need for synthetic vehicular mobility datasets**
  We first review various means of collecting and representing the vehicular mobility data available today. In this process, we highlight the shortcomings of real-world vehicular mobility datasets and discuss how such attempts are less attractive for design and evaluation of network protocols. We present the traditional procedures and tools involved in the generation of synthetic vehicular mobility dataset along with their timeline of evolution. We then review various synthetically generated mobility datasets with varying level of realism in their macroscopic data and microscopic simulation. We see that most of these synthetically generated datasets either cover a small geographical area, or

span over a short period of time or involve mobility over major arteries. Hence these datasets fail to provide a sufficient platform for large-scale protocol evaluation and may possibly lead to unreliable results or inconclusive studies.

Proposal and results : We introduce a large-scale synthetic vehicular mobility dataset which is the result of using highly realistic macroscopic data with detailed microscopic simulation. Apart from the realism and detail of the source data and models, several additional procedures need to be performed to achieve a realistic dataset that matches with the real-world traffic in Koln, Germany. To the best of our knowledge, the dataset we present constitutes the most complete vehicular mobility dataset to date, that extends over an area of 400 km^2 during 24 hours providing mobility information with high time granularity. In particular, the SUMO configuration data is publicly available, which makes our dataset usable with bidirectional simulation tools, such as TraNS [PRL+08], Veins [VEI12], or iTetris [iTe11], and allows to dynamically modify the road traffic as a consequence of messages exchanged in the vehicular network.

- **Spatiotemporal implication of vehicular mobility on cellular networks and autonomous network architectures**

  The single most critical aspect that makes vehicular environments especially challenging is its unique mobile nature. Also, it is to be said that works on the access network have mainly focused on infrastructures dedicated to vehicular access via Dedicated Short-Range Communication (DSRC) rather than through traditional cellular RANs. Hence, there is a need to understand the fundamentals of pervasive vehicular access in cellular RANs, as it has been been done in the case of pedestrian users for the past thirty years or so [HR86, BZ97, LGFA]. Due to the lack of detailed vehicular mobility dataset most studies in autonomous networks are either limited by time, small-scale or rely on simplistic assumptions on the vehicle mobility, such as unidimensional mobility and atomic contacts among vehicles. Hence there is a need for a large-scale mobility characterization to understand the features of vehicular movement patterns that can help in determining the potential vehicle-to-vehicle data transfers, in estimating future communication opportunities among moving vehicles, and also in planning the deployment of roadside infrastructure.

  Proposal and results : Leveraging a real cellular infrastructure deployment information in the city of Koln, we mimic the cellular coverage with Voronoi tessellation. We learned how the network load varies with time and space in an urban environment, and direct assumption of network load derived from population density does not mimic the vehicular access behavior. Also inter-cell vehicular mobility dynamics follows specific

patterns with space and time with traffic flows traversing in/out of the city. Studying these mobility patterns uncovered that the user arrival into a cellular cell does not necessarily follow a Poisson distribution as assumed in most cellular network protocol evaluations. Due to high mobility dynamics, vehicular access exhibits remarkably low cell resident times, hinting the use of umbrella cells to ease the data traffic demand and lowering the need for dedicated resources used for control message. We also evaluate the impact of realistic vehicular mobility on the connectivity of large- and small-scale autonomous networks of vehicles. We employ various networking metrics and compare the Koln dataset with other freely available datasets with an autonomous network connectivity perspective. The results bring us to conjecture that evaluating network protocols or architectures through low-detail or spatially- and temporally-limited mobility traces is a risky practice, that can lead to over-optimistic performance results.

- **Need for lightweight online macroscopic vehicular mobility prediction and analysis**

  Research efforts in the past have established that the human mobility is often regular and can be predicted. However most of these studies are based on datasets concentrating on small group of population or involving mobility data of public transports which are often regular and less challenging. Moreover these prediction attempts are based on large historical data which are hard to obtain and have important processing and memory requirements. In addition, growing privacy concerns of the users makes these approaches practically less encouraging.

  Proposal and results : We proposed a lightweight prediction approach which is online in nature to predict the future macroscopic traffic flows across the crossroads and cellular cells. This approach profiles users based on their similar mobility behaviors, tracking flows of users rather than individual users, which address the privacy, processing and memory requirements. This study unveils the possibility of mobility prediction with movement pattern knowledge over a short period of past time. The parameters affecting the mobility prediction play an important role in prediction accuracy, and hence require attention when incorporating such mechanisms in real-world systems. The spatiotemporal analysis uncovers that the prediction at the outskirts of the city often requires more historical knowledge. But in the case of city center, the vehicular mobility can be predicted with around 90% accuracy with shorter historical knowledge. Such a finding hints at possibilities for intelligent resource management in lightweight cellular base stations that are introduced to meet the capacity requirement in recent developments of the cellular technology.

6.2 Perspectives

Though this study addresses several of the requirement that are significant for the evaluation of network solutions, there are still many opportunities to extend this work.

- **Realism of state-of-art tools**
  The road to the generation of ultimate synthetic vehicular mobility datasets is still long. More precisely, aspects that are still to be addressed concern at a time the scalability of microscopic simulation and the increased detail of the same. The generation of Koln vehicular mobility dataset involved real world road map that was repaired or simplified manually to integrate well with SUMO simulator. Since manual procedures are time consuming and lacks perfection, automated procedures to fix the road network representation and traffic rules are needed to further increase the level of realism. In terms of scalability, it would be desirable to obtain datasets that cover whole district/states, and, why not, even countries. Moreover, the macroscopic traffic description remains an open problem. Apart from isolated initiatives, such as the TAPASCologne one, realistic road traffic demand information on vast regions is hard to retrieve. To that end, the growing trend of government bodies and local administrations making urban-related data publicly available can provide a significative breakthrough. Open access data initiatives such as Vancouver's VanMap, OpenBaltimore, London DataStore and ParisData, just to cite a few significant examples, could indeed lead to the public disclosure of road counter collections or traffic flow datasets. Such information would, in turn, grant the possibility to extract new macroscopic O/D matrices, as well as validate traffic assignments.
- **Validating realism of datasets through networking applications**
  The Macroscopic and microscopic analysis with Voronoi tessellation establishes a benchmark to the nature of mobility dynamics in Koln urban environment in space and time. More tests in different urban environments are needed to verify the general validity of these results. Also additional network level simulation in autonomous network is required to demonstrate the level of realism of different communication protocols.
- **Validating the prediction approach**
  The proposed prediction approach could be incorporated into a cellular application like dynamic resource management or efficient handoff management to demonstrate the effectiveness of our lightweight approach. Also the prediction result could be quantified against a traditional method that relies on large traffic dynamic knowledge. Further, different prediction algorithms could be involved to check the variability of prediction error against the Markovian and Sequence Matching models.

7 List of publications

- **Book Chapter :**
  1. S. Uppoor, M. Fiore, J. Härri. *Synthetic mobility traces for vehicular networking.* Vehicular Networks, H. Labiod and A.-L. Beylot (Editors), pp.209-245, John Wiley & Sons, 2013, ISBN: 9781848214897.

- **Journals :**
  1. S. Uppoor, M. Fiore. *Characterizing pervasive vehicular access to the cellular RAN infrastructure: an urban case study.* Submitted.
  2. 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.
  3. S. Uppoor, M. Fiore. *Vehicular mobility in large-scale urban environments.* ACM Mobile Computing and Communications Review, Vol.15, No.4, 2011.

- **International Conferences :**
  1. S. Uppoor, M. Fiore. *Insights on metropolitan-scale vehicular mobility from a networking perspective.*Invited paper, ACM HotPlanet 2012, Low Wood Bay, UK.
  2. S. Uppoor, M. Fiore. *Large-scale Urban Vehicular Mobility for Networking Research.* IEEE VNC 2011, Amsterdam, Netherlands.
  3. S. Uppoor, M. Fiore. *Vehicular mobility in large-scale urban environments.* MobiCom 2011 Poster Session, Las Vegas, USA.

- **National Conferences and Workshops :**
  1. Sandesh Uppoor, Marco Fiore. *A large-scale vehicular mobility dataset of the Cologne urban area .* Algotel 2012, Hérault, France.

2. S. Uppoor, M. Fiore. *A large-scale urban vehicular mobility trace for network research.* JNCT 2011, Colmar, France.
3. S. Uppoor, M. Fiore. *Urban road traffic modeling and impact on vehicular networking.* COST IC1004 Scientific meeting and Workshop, Lyon, France.

- **Demo :**
  1. L. Roullet, J. Härri, N. Nikaein, S. Uppoor, M. Fiore. *Scalable mobility over cellular networks.* Bell Labs Open Days 2012, Bell Labs Research Center, Villarceaux France.

Bibliography

[AIM12] AIMSUN. Aimsun. http://www.aimsun.com/wp/, 2012.

[AL] Alcatel-Lucent. Beyond the base station router. *Alcatel-Lucent technical note*. http : //innovationdays.alcatel ?lucent.com/2008/documents/Beyond%20BSR.pdf.

[AMSS11] Akinori Asahara, Kishiko Maruyama, Akiko Sato, and Kouichi Seto. Pedestrian-movement prediction based on mixed markov-chain model. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 25–33. ACM, 2011.

[AS02] Daniel Ashbrook and Thad Starner. Learning significant locations and predicting user movement with gps. In *Wearable Computers, 2002.(ISWC 2002). Proceedings. Sixth International Symposium on*, pages 101–108. IEEE, 2002.

[AS07] Sherif Akoush and Ahmed Sameh. Movement prediction using bayesian learning for neural networks. In *ICSNC'07*, pages –1–1, 2007.

[AUT12] AUTOBAHN. Traffic information system, ministry of transport of nordrhein-westfalen. http://www.autobahn.nrw.de/, 2012.

[BD99] Amiya Bhattacharya and Sajal K Das. Lezi-update: an information-theoretic approach to track mobile users in pcs networks. In *Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*, pages 1–12. ACM, 1999.

[BG09] C. Barberis and Malnati G. Large-scale simulation of v2v environments. In *IEEE ICUMT*, St. Petersburg, Russia, 2009.

[BGJL06] J. Burgess, B. Gallagher, D. Jensen, and B.N. Levine. Maxprop: Routing for vehicle-based disruption-tolerant networks. In *IEEE INFOCOM*, Barcelona, Spain, 2006.

[BHM+06] V Bychkovsky, B Hull, A Miu, H Balakrishnan, and S Madden. A measurement study of vehicular internet access using in situ wi-fi networks. In *Proceedings of the 12th annual international conference on Mobile computing and networking*,

pages 50–61. ACM, 2006.

[BK09] F. Bai and B. Krishnamachari. Spatio-temporal variations of vehicle traffic in vanets: facts and implications. In *ACM VANET*, Beijing, PRC, 2009.

[BLJL10] A Bohm, K Lidstrom, Magnus Jonsson, and Tony Larsson. Evaluating calm m5-based vehicle-to-vehicle communication in various road settings through field trials. In *Local Computer Networks (LCN), 2010 IEEE 35th Conference on*, pages 613–620. IEEE, 2010.

[BLP08] R. Baumann, F. Legendre, and Sommer P. Generic mobility simulation framework (gmsf). In *ACM MobilityModels*, Hong Kong, PRC, 2008.

[Bra11] R. Neil Braithwaite. Improving data throughput for cell-edge users in a lte network using up-link harq relays. In *VTC Fall*, pages 1–5. IEEE, 2011.

[Bre01] L. Breisemeister. *Group Membership and communication in highly mobile ad hoc networks*. PhD thesis, Tecnhical University of Berlin, 2001.

[BSH03] F. Bai, N. Sadagopan, and A. Helmy. The IMPORTANT framework for analyzing the Impact of Mobility on Performance Of RouTing protocols for Adhoc NeTworks. *Elsevier Ad Hoc Networks*, 1:383–403, 2003.

[BSK10] Fan Bai, Daniel D. Stancil, and Hariharan Krishnan. Toward understanding characteristics of dedicated short range communications (dsrc) from a perspective of vehicular network engineers. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, MobiCom '10, pages 329–340, New York, NY, USA, 2010. ACM.

[Bur12] U.S. Census Bureau. Topologically integrated geographic encoding and referencing (tiger) system. http://www.census.gov/geo/www/tiger, 2012.

[BZ97] François Baccelli and Serguei Zuyev. Stochastic geometry models of mobile communication networks. 1997.

[Cab11] CabSpotting. http://cabspotting.org/api, 2011.

[CaDFJ08] H. Conceição, L. Damas, M. Ferreira, and Barros J. Large-scale simulation of v2v environments. In *ACM SAC*, Fortaleza, Brazil, 2008.

[CB04] Joe Capka and Raouf Boutaba. Mobility prediction in wireless networks using neural networks. In *MMNS'04*, pages 320–334, 2004.

[CBN03] Nurhan Cetin, Adrian Burri, and Kai Nagel. A large-scale agent-based traffic microsimulation based on queue model. In *Swiss Transport Research Conference (STRC), MONTE VERITA, CH*, pages 23–41, 2003.

BIBLIOGRAPHY

[Chi82] *The Chicago Manual of Style*, pages 400–401. University of Chicago Press, thirteenth edition, 1982.

[CHS+07] Lin Cheng, B. E. Henty, D. D. Stancil, Fan Bai, and P. Mudalige. Mobile vehicle-to-vehicle narrow-band channel measurement and characterization of the 5.9 ghz dedicated short range communication (dsrc) frequency band. *IEEE J.Sel. A. Commun.*, 25(8):1501–1516, October 2007.

[Cis13] Cisco. Global mobile data traffic forecast update 2012 – 2017. White Paper, 2013.

[Civ06] A. Civilis. Prediction of crossroad passing using artificial neural networks. In *Databases and Information Systems, 2006 7th International Baltic Conference on*, pages 229 –234, 0-0 2006.

[COR12] CORSIM. Tsis-corsim. http://mctrans.ce.ufl.edu/featured/tsis/, 2012.

[Cor13] SW Corbett. Traffic calming project evaluations. http://www.portlandoregon.gov/transportation/article/85277, 2013.

[CSZ11] Ho Ting Cheng, Hangguan Shan, and Weihua Zhuang. Infotainment and road safety service support in vehicular networking: From a communication perspective. *Mechanical Systems and Signal Processing*, 25(6):2020–2038, 2011.

[CZS98] Jonathan Chan, S Zhou, and Aruna Seneviratne. A qos adaptive mobility prediction scheme for wireless networks. In *Global Telecommunications Conference, 1998. GLOBECOM 1998. The Bridge to Global Integration. IEEE*, volume 3, pages 1414–1419. IEEE, 1998.

[Dav00] V. Davies. Evaluating mobility models within an ad hoc network. Master's thesis, Colorado School of Mines, 2000.

[DK04] Mukund Deshpande and George Karypis. Selective markov models for predicting web page accesses. *ACM Transactions on Internet Technology (TOIT)*, 4(2):163–184, 2004.

[dlVP79] Charles Louis Xavier Joseph de la Vallée Poussin, 1879. A strong form of the prime number theorem, 19th century.

[DMW+11] Aleksandar Damnjanovic, Juan Montojo, Yongbin Wei, Tingfang Ji, Tao Luo, Madhavan Vajapeyam, Taesang Yoo, Osok Song, and Durga Malladi. A survey on 3gpp heterogeneous networks. *Wireless Communications, IEEE*, 18(3):10–21, 2011.

[DPPW10] M. Doering, T. Pögel, W.-B. Pöttner, and L. Wolf. A new mobility trace for realistic large-scale simulation of bus-based dtns. In *ACM CHANTS*, Chicago,

IL, USA, 2010.

[EB96] M. Ehling and W. Bihler. Zeit im blickfeld. ergebnisse einer repräsentativen zeitbudgeterhebung. *Schriftenreihe des Bundesministeriums für Familie, Senioren, Frauen und Jugend*, 121:237–274, 1996.

[Ert12] Ertico. Geographic data files. http://www.ertico.com/gdf-geographic-data-files, 2012.

[ESE+01] Fazli Erbas, Jan Steuer, Dirk Eggesieker, Kyandoghere Kyamakya, and K Jobinann. A regular path recognition method and prediction of user movements in wireless networks. In *Vehicular Technology Conference, 2001. VTC 2001 Fall. IEEE VTS 54th*, volume 4, pages 2672–2676. IEEE, 2001.

[ETS13] ETSI. Etsi. http://www.etsi.org/technologies-clusters/technologies/intelligent-transport, 2013.

[EUR12] EURECOM. Openairinterface. http://www.openairinterface.org/., 2012.

[FAS12] FASTTRANS. Fasttrans. http://www.lanl.gov/programs/nisac/fasttrans.shtml, 2012.

[FCCP12] M. Fiore, C. Casetti, C.-F. Chiasserini, and P. Papadimitratos. Discovery and verification of neighbor positions in mobile ad hoc networks. *IEEE Transactions on Mobile Computing*, to appear in 2012.

[FCFO09] M. Ferreira, H. Concei, R. Fernandes, and Tonguz O.K. Stereoscopic aerial photography: an alternative to model-based urban mobility approaches. In *ACM VANET*, Beijing, PRC, 2009.

[FG12] FST-GSM. Mobilfunk in köln. http://www.fst-gsm.de/, 2012.

[FH08] M. Fiore and J. Härri. The Networking Shape of Vehicular Mobility. In *ACM MobiHoc*, Hong Kong, PRC, 2008.

[FHFB07] M. Fiore, J. Härri, F. Filali, and C. Bonnet. Vehicular mobility simulation for vanets. In *SCS/IEEE ANSS*, Norfolk, VA, USA, 2007.

[FMR06] R. Fracchia, M. Meo, and D. Rossi. Vanets: To beacon or not to beacon. In *IEEE Globecom Autonet*, San Francisco, CA, USA, 2006.

[Gaw98] C. Gawron. An iterative algorithm to determine the dynamic user equilibrium in a traffic simulation model. *International Journal of Modern Physics C*, 9(3):393–407, 1998.

[GC07] Karthik Gopalratnam and Diane J Cook. Online sequential prediction via incremental parsing: The active lezi algorithm. *Intelligent Systems, IEEE*, 22(1):52–

58, 2007.

[Gip81] P.G. Gipps. A behavioural car following model for computer simulation. *Transportation Research B*, 15:105–111, 1981.

[GKdPC10] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. Show me how you move and i will tell you who you are. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Security and Privacy in GIS and LBS*, pages 34–41. ACM, 2010.

[GR02] H. Mühlhans G. Rindsfüser, J. Ansorge. *SimVV Mobilität verstehen und lenken – zu einer integrierten quantitativen Gesamtsicht und Mikrosimulation von Verkehr*, chapter Aktivitätenvorhaben. Final report, Ministry of School, Science and Research of Nordrhein-Westfalen, 2002.

[Har09] J Harri. *Vehicular Mobility Modeling for VANET*, pages 107–156. John Wiley & Sons, Ltd, 2009.

[HFC09] J. Härri, F. Filali, and Bonnet C. Mobility models for vehicular ad hoc networks: a survey and taxonomy. *IEEE Communications Surveys and Tutorials*, 11(4):19–41, 2009.

[HFFB11] J. Härri, M. Fiore, F. Filali, and C. Bonnet. Vehicular mobility simulation with vanetmobisim. *Transactions of The Society for Modeling and Simulation International*, 87(4), 2011.

[HLL+07] H.-Y. Huang, P.-E. Luo, M. Li, D. Li, X. Li, W. Shu, and Wu. M.-Y. Performance evaluation of suvnet with real-time traffic data. *IEEE Transactions on Vehicular Technology*, 56(6):3381–3396, 2007.

[HR86] Daehyoung Hong and Stephen S Rappaport. Traffic model and performance analysis for cellular mobile radio telephone systems with prioritized and non-prioritized handoff procedures. *Vehicular Technology, IEEE Transactions on*, 35(3):77–92, 1986.

[HW04] G. Hertkorn and P. Wagner. The application of microscopic activity based travel demand modelling in large scale simulations. In *World Conference on Transport Research*, Istanbul, Turkey, 2004.

[iTe11] iTetris. An integrated wireless and traffic platform for real-time road traffic management solutions. http://ict-itetris.eu, 2011.

[IUT13] IUT. Ict facts and figures 2013. http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2013.pdf, 2013.

[JBL05] S. Jaap, M. Bechler, and Wolf L. Evaluation of routing protocols for vehicular

ad hoc networks in city traffic scenarios. In *IEEE ITSC*, Vienna, Austria, 2005.

[JD08] Daniel Jiang and Luca Delgrossi. Ieee 802.11 p: Towards an international standard for wireless access in vehicular environments. In *Vehicular Technology Conference, 2008. VTC Spring 2008. IEEE*, pages 2036–2040. IEEE, 2008.

[JDS12] Stefan Joerer, Falko Dressler, and Christoph Sommer. Comparing apples and oranges?: trends in ivc simulations. In *Proceedings of the ninth ACM international workshop on Vehicular inter-networking, systems, and applications*, VANET '12, pages 27–32, New York, NY, USA, 2012. ACM.

[JHP+03] J.G. Jetcheva, Y.-C. Hu, S. PalChaudhuri, A.K. Saha, and D.B. Johnson. Design and evaluation of a metropolitan area multitier wireless ad hoc network architecture. In *IEEE WMCSA*, Monterey, CA, USA, 2003.

[JLSZ08] Hoyoung Jeung, Qing Liu, Heng Tao Shen, and Xiaofang Zhou. A hybrid prediction model for moving objects. In *Data Engineering, 2008. ICDE 2008. IEEE 24th International Conference on*, pages 70–79. IEEE, 2008.

[JOS12] JOSM. Java openstreetmap editor,. http://josm.openstreetmap.de/, 2012.

[JYZJ10] Hoyoung Jeung, Man Lung Yiu, Xiaofang Zhou, and Christian S Jensen. Path prediction and predictive range querying in road network databases. *The VLDB Journal*, 19(4):585–602, 2010.

[KHRW02] D. Krajzewicz, G. Hertkorn, C. Rossel, and P. Wagner. Sumo (simulation of urban mobility): An open-source traffic simulation. In *SCS MESM*, Sharjah, United Arab Emirates, 2002.

[KK90a] K.S.Narendra and K.Parthsarathy. Identification and control of dynamical system using neural networks. *IEENN*, 1(1):4–27, 1990.

[KK90b] K.S.Narendra and K.Parthsarathy. Identification and control of dynamical system using neural networks. In *Infocom*, pages 4–27, New York, 1990. Wiley, New York.

[KML07] F. Karnadi, Z. Mo, and K.-C. Lan. Rapid generation of realistic mobility models for vanet. In *IEEE WCNC*, Hong Kong, 2007.

[Knu84] Donald E. Knuth. *The TEXbook*. Addison-Wesley, 1984.

[KOL11] KOLN. Vehicular mobility trace of the city of cologne, germany. http://kolntrace.project.citi-lab.fr/, 2011.

[Kra98] S. Krauss. *Microscopic Modeling of Traffic Flow: Investigation of Collision Free Vehicle Dynamics*. PhD thesis, Universität zu Köln, 1998.

[Kra09] D. Krajzewicz. Kombination von taktischen und strategischen einflüssen in einer mikroskopischen verkehrsflusssimulation. pages 104–115, 2009.

[Kru08] John Krumm. A markov model for driver turn prediction. SAE Technical Paper, 2008.

[KWG97] S. Krauss, P. Wagner, and C. Gawron. Metastable states in a microscopic model of traffic flow. *Physical Review E*, 55(304):55–97, 1997.

[KZ97] F. Khan and D. Zeghlache. Effect of cell residence time distribution on the performance of cellular mobile networks. In *Vehicular Technology Conference, 1997, IEEE 47th*, volume 2, pages 949–953, 1997.

[Lam86] Leslie Lamport. *LaTeX: A Document Preparation System*. Addison-Wesley, 1986.

[Las06] F Lassabe. Predictive mobility models based on kth markov models. pages 303 – 306. ACS/IEEE International Conference on Pervasive Services, June 2006.

[Lev44] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly Journal of Applied Mathmatics*, II(2):164–168, 1944.

[LGFA] Xingqin Lin, R Ganti, Philip Fleming, and J Andrews. Towards understanding the fundamentals of mobility in cellular networks.

[LHT+03] C. Lochert, H. Hartenstein, J. Tian, H. Fuessler, D. Hermann, and M. Mauve. A routing strategy for vehicular ad hoc networks in city environments. In *IEEE IV*, Columbus, OH, USA, 2003.

[LM96] George Liu and Gerald Maguire, Jr. A class of mobile motion prediction algorithms for wireless mobile computing and communication. *Mob. Netw. Appl.*, 1(2):113–121, oct 1996.

[LWY+06] M. Li, M.-Y. Wu, Y. Ying, J. Cao, L. Huang, Q. Deng, X. Lin, C. Jiang, W. Tong, Y. Gui, A. Zhou, X. Wu, and S. Jiang. Shanghaigrid: an information service grid. In *Concurrency and Computation: Practice and Experience*. John Wiley & Sons, 2006.

[MAG07] WIRED MAGAZINE. The original futurama. *WIRED MAGAZINE ONLINE EXTRAS*, 15(02), November 2007.

[Mar63] D. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.

[MAT12] MATSim. Matsim. http://matsim.org, 2012.

[Mei11] MeihuiSoftware. http://www.meihuichina.com, 2011.

[M.F09] M.Fiore. *Vehicular Networks from Theory to Practice*, chapter 12 – Vehicular Mobility Models. Chapman & Hall/CRC, 2009.

[NS92] K. Nagel and M. Schreckenberg. A cellular automaton model for freeway traffic. *Journal de Physique I*, 2(12):2221–2229, 1992.

[NS312] NS3. Network simulator 3. http://www.nsnam.org/., 2012.

[NTD+06] A. Nandan, S. Tewari, S. Das, M. Gerla, and L. Kleinrock. Adtorrent: Delivering location cognizant advertisements to car networks. In *WONS*, Les Ménuires, France, 2006.

[NWWS98] K. Nagel, D.E. Wolf, P. Wagner, and P. Simon. Two-lane traffic rules for cellular automata: a systematic approach. *Physical Review E*, 58:1425–1437, 1998.

[OMN12] OMNET++. Omnet++. http://www.omnetpp.org/., 2012.

[Ope12] OpenStreetMap. The free wiki world map. http://www.openstreetmap.org, 2012.

[Osm12] Osmosis. Osmosis. http://wiki.openstreetmap.org/wiki/Osmosis/, 2012.

[PAR12] PARAMICS. Paramics. http://paramics.com/, 2012.

[Pat88a] Oren Patashnik. Designing bibtex styles. The part of BibTeX's documentation that's not meant for general users, January 1988.

[Pat88b] Oren Patashnik. Using BibTeX. Documentation for general BibTeX users, January 1988.

[PDP11] Y. Pigne, G. Danoy, and Bouvry P. A vehicular mobility model based on real traffic counting data. In *Nets4Cars*, Oberpfaffenhofen, Germany, 2011.

[PKSG11] Jeongyeup Paek, Kyu-Han Kim, Jatinder P Singh, and Ramesh Govindan. Energy-efficient positioning for smartphones using cell-id sequence matching. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*, pages 293–306. ACM, 2011.

[PLFK03] DonaldJ. Patterson, Lin Liao, Dieter Fox, and Henry Kautz. Inferring high-level behavior from low-level sensors. In AnindK. Dey, Albrecht Schmidt, and JosephF. McCarthy, editors, *UbiComp 2003: Ubiquitous Computing*, volume 2864 of *Lecture Notes in Computer Science*, pages 73–89. Springer Berlin Heidelberg, 2003.

[PRL+08] Michal Piorkowski, Maxim Raya, A Lezama Lugo, Panagiotis Papadimitratos, Matthias Grossglauser, and J-P Hubaux. Trans: realistic joint traffic and network simulator for vanets. *ACM SIGMOBILE Mobile Computing and Communications Review*, 12(1):31–33, 2008.

[PSBD11] U. Paul, A.P. Subramanian, M.M. Buddhikot, and S.R. Das. Understanding

traffic dynamics in cellular data networks. In *INFOCOM, 2011 Proceedings IEEE*, pages 882–890, Shanghai, China, 2011.

[RCV+03] B. Raney, N. Cetin, A. Völlmy, M. Vrtic, K. Axhausen, and K. Nagel. An agent-based microsimulation model of swiss travel: First results. *Networks and Spatial Economics*, 3(1):23–41, 2003.

[RKA10] Sreenath Ramanath, Veeraruna Kavitha, and Eitan Altman. Impact of mobility on call block, call drops and optimal cell size in small cell networks. In *Personal, Indoor and Mobile Radio Communications Workshops (PIMRC Workshops), 2010 IEEE 21st International Symposium on*, pages 157–162. IEEE, 2010.

[RML02] M. Rudack, M. Meincke, and M. Lott. On the dynamics of ad hoc networks for inter-vehicle communications (ivc). In *ICWN*, Las Vegas, NV, USA, 2002.

[SB07] Michael H Sun and Douglas M Blough. Mobility prediction using future knowledge. In *Proceedings of the 10th ACM Symposium on Modeling, analysis, and simulation of wireless and mobile systems*, pages 235–239. ACM, 2007.

[SBrA05] S.Karpinski, E.M. Belding-royer, and K.C. Almeroth. Living on the edge: The distribution of flows across mobile nodes in large wireless networks. Technical Report., 2005.

[Sch07] R Schougaard, Kari. Vehicular mobility prediction by bayesian networks. Technical report, 2007.

[SDK+06] L Song, U Deshpande, U.C Kozat, D Kotz, and R Jain. Predictability of wlan mobility and its effects on bandwidth provisioning. INFOCOM, 2006.

[SGD11] C. Sommer, R. German, and F. Dressler. Bidirectionally coupled network and road traffic simulation for improved ivc analysis. *Mobile Computing, IEEE Transactions on*, 10(1):3–15, 2011.

[She13] ITS Standards Fact Sheets. Ieee 1609 – family of standards for wireless access in vehicular environments (wave). http://www.standards.its.dot.gov/Factsheets/Factsheet/80, 2013.

[SIm13] SImTD. Simtd. http://www.simtd.de, 2013.

[SJ04] A.K. Saha and D.B. Johnson. Modeling mobility for vehicular ad hoc networks. In *ACM VANET*, Philadelphia, PA, USA, 2004.

[SKJH04a] Libo Song, David Kotz, Ravi Jain, and Xiaoning He. Evaluating location predictors with extensive wi-fi mobility data. In *INFOCOM 2004. Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies*, volume 2, pages 1414–1424. IEEE, 2004.

[SKJH04b] Libo Song, David Kotz, Ravi Jain, and Xiaoning He. Evaluating location predictors with extensive wi-fi mobility data. In *INFOCOM 2004. Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies*, volume 2, pages 1414–1424. IEEE, 2004.

[SLL+08] M. Sede, X. Li, D. Li, M.-Y. Wu, M. Li, and W. Shu. Routing in large-scale buses ad hoc networks. In *IEEE WCNC*, Las Vegas, NV, USA, 2008.

[SMHW92] I. Seskar, S. Marie, J. Holtzman, and J. Wasserman. Rate of location area updates in cellular systems. In *IEEE VTC*, Denver, CO, USA, 1992.

[SUM12] SUMO. Simulation of urban mobility. http://sumo.sourceforge.net/, 2012.

[SW79] William Strunk, Jr. and E. B. White. *The Elements of Style*. Macmillan, third edition, 1979.

[SWA12] SWANS. Scalable wireless ad-hoc network simulator. http://www.jist.ece.cornell.edu/., 2012.

[TH02] M. Treiber and D. Helbing. Realistische mikrosimulation von strassenverkehr mit einem einfachen modell. In *ASIM*, Rostock, Germany, 2002.

[THB+02] J. Tian, J. Haehner, C. Becker, I. Stepanov, and K. Rothermel. Graph-based mobility model for mobile ad hoc network simulation. In *SCS ANSS*, San Diego, CA, USA, 2002.

[THH00] M. Treiber, A. Hennecke, and D. Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2):1805–1824, 2000.

[THKH11] G.S. Thakur, P. Hui, H. Ketabdar, and A. Helmy. Towards realistic vehicular network modeling using planet-scale public webcams. In *Computing Research Repository (CoRR)*, 2011.

[Tom10] TomTom. *How TomTom's HD TrafficTM and IQ RoutesTM data provides the very best routing*. White paper, 2010.

[TRA12] TRANSIMS. Transportation analysis and simulation system. http://code.google.com/p/transims/, 2012.

[UF11] S. Uppoor and M. Fiore. Large-scale urban vehicular mobility for networking research. In *IEEE VNC*, Amsterdam, The Netherlands, 2011.

[UF12] Sandesh Uppoor and Marco Fiore. Insights on metropolitan-scale vehicular mobility from a networking perspective. In *Proceedings of the 4th ACM international workshop on Hot topics in planet-scale measurement*, HotPlanet '12, pages 39–44, New York, NY, USA, 2012. ACM.

[Van12] VanetMobiSim. The vehicular mobility simulator. http://sourceforge.net/projects/vanetmobisim/, 2012.

[VBT11a] W. Viriyasitavat, F. Bai, and O.K. Tonguz. Dynamics of Network Connectivity in Urban Vehicular Networks. *IEEE Journal on Selected Areas in Communications*, 29(3):515–533, 2011.

[VBT11b] Wantanee Viriyasitavat, Fan Bai, and Ozan K Tonguz. Dynamics of network connectivity in urban vehicular networks. *Selected Areas in Communications, IEEE Journal on*, 29(3):515–533, 2011.

[VEI12] VEINS. Veins. http://www.veins.car2x.org/., 2012.

[VIS12a] VISSIM. Vissim. http://www.ptv-vision.com/en-uk/products/vision-traffic-suite/ptv-vissim/overview/, 2012.

[VIS12b] VISUM. Visum. http://www.ptv-vision.com/en-uk/products/vision-traffic-suite/ptv-vissim/overview/., 2012.

[VK96] Jeffrey Scott Vitter and P Krishnan. Optimal prefetching via data compression. *Journal of the ACM (JACM)*, 43(5):771–793, 1996.

[VKV02] B. P. Vijay, Kumar, and P. Venkataram. Prediction-based location management using multilayer neural networks. *Journal of Indian institute of science*, pages 7–21, 2002.

[vL79] Mary-Claire van Leunen. *A Handbook for Scholars*. Knopf, 1979.

[VSI12] VSIMRTI. Vsimrti. http://www.dcaiti.tu-berlin.de/research/simulation/., 2012.

[VW06] C. Varschen and P. Wagner. Mikroskopische modellierung der personenverkehrsnachfrage auf basis von zeitverwendungstagebüchern. *Stadt Region Land*, 81:63–69, 2006.

[Wie74] R Wiedemann. Simulation des straenverkehrsflusses, schriftenreihe des ifv, 1974.

[WK99] P. Wagner and Nagel K. Microscopic modeling of travel-deman: Approaching the home-to-work problem. In *Transportation Research Board Annual Meeting*, Washington, DC, USA, 1999.

[WLY+11] Weetit Wanalertlak, Ben Lee, Chansu Yu, Myungchul Kim, Seung-Min Park, and Won-Tae Kim. Behavior-based mobility prediction for seamless handoffs in mobile wireless networks. *Wireless Networks*, 17(3):645–658, 2011.

[WPR+08] A Wegener, M Piorkowski, M Raya, H Hellbrck, S Fischer, and J-P Hubaux. Traci: an interface for coupling road traffic and network simulators. In *Proceedings of the 11th communications and networking simulation symposium*, pages

155–163. ACM, 2008.

[XLYL12] G Xue, Y Luo, J Yu, and M Li. A novel vehicular location prediction based on mobility patterns for routing in urban vanet. *EURASIP Journal on Wireless Communications and Networking*, 2012(1):1–14, 2012.

[YKUM05] Gökhan Yavaş, Dimitrios Katsaros, Özgür Ulusoy, and Yannis Manolopoulos. A data mining approach for location prediction in mobile environments. *Data & Knowledge Engineering*, 54(2):121–146, 2005.

[YZXS11] J. Yuan, Y. Zheng, X. Xie, and G. Sun. Driving with knowledge from the physical world. In *ACM SIGKDD*, San Diego, CA, USA, 2011.

[ZAKB10] L. Zhang, J. Ahn, B. Krishnamachari, and F. Bai. Optimizing content dissemination in heterogeneous vehicular networks. In *PATH/Tsinghua Workshop on ITS*, Berkeley, CA, USA, 2010.

[ZM93] R. W. Zurek and L. J. Martin. Interannual variability of planet-encircling dust activity on Mars. *J. Geophys. Res.*, 98(E2):3247–3259, 1993.