Service-Oriented Information-Centric Vehicular Ad-hoc Networks

by

Felipe Modesto

Thesis submitted to the
Faculty of Graduate and Postdoctoral Studies
In partial fulfillment of the requirements
For the Ph.D. degree in
Computer Science

School of Information Technology and Engineering
Faculty of Engineering
University of Ottawa

Abstract

With Vehicular mobile communication becoming a daily requirement and an ever increasing number of services being available to passengers, it is clear that vehicular networks efficient communication systems. VANETs, one of the most significant trends in ad-hoc networking, has much to gain from improved content delivery and one of the leading contenders for mobile networks is the Information-Centric networking approach. Its peculiarities define the Vehicular Environment requires specialized solutions, tailored for highly mobile environments. The main contribution of this thesis is the introduction of a novel architecture and components. We perform extensively discuss Information-Centric Vehicular Ad-hoc Networks. Additionally, we perform an in-depth analysis of bus-based transit systems into VANETs not only as participating members but as service providers and official agents including roles and potential challenges. We perform statistical analysis and analyze world data to denote the intrinsic potential of public transit systems. From the discussions presented, we introduce a novel service-based system architecture for Information-Centric Networking named SEVeN. The proposed model is designed to enable service exchange and service management in highly competitive vehicular adhoc networks. The proposed SEVeN architecture includes the introduction of a novel purpose-defined naming policy and service sub-layer as well as a service prioritization policy named LBD. We also discuss the current state of ICN caching in VANET, existing issues faced by vehicular networks and potential approaches based on intermediate cache coordination that can be taken to mitigate existing shortcommings. We perform a series of simulations and analyze the efficiency of popular caching in various network configurations to denote current shortcomings. From this discussion, we propose a cache content insertion policies, UG-Cache and MG-Cache, for ICN-VANETs. In these cache policies, cache insertion decisions are made based on recommendations from content sender dependent on request frequency and cache distance. We also introduce a caching policy based on collaborative observation of locality in request frequency, designed to allow vehicles to preemptively distribute and store in a reserved portion of the cache based on the cooperative observation of requests with provider-based location correlation. All novel elements proposed by this thesis are discussed, described, evaluated within the chapters of this thesis.

List of Publications

Referred Conference and Journal Papers

1. F. Modesto and A. Boukerche, "Towards Integrating Public Transit Bus Systems into Urban and Intelligent Vehicular Networks", *2019 IEEE Wireless Communications and Networking Conference (WCNC)*, pp. xxx–xxx, 2019. doi: `10.1109/WCNC.2019.XXXXXXX`

2. F. Modesto and A. Boukerche, "Social Pre-caching for Location-dependent Requests in Vehicular Information-Centric Ad-hoc Networks", *2018 IEEE Symposium on Computers and Communications (ISCC)*, pp. 426–431, 2018. doi: `10.1109/ISCC.2018.8538613`

3. F. Modesto and A. Boukerche, "Utility-Gradient Implicit Cache Coordination Policy for Information-Centric Ad-hoc Vehicular Networks", *2017 IEEE 42nd Conference on Local Computer Networks (LCN)*, pp. 010–017, 2017. doi: `10.1109/LCN.2017.61`

4. F. Modesto and A. Boukerche, "SEVeN: A novel service-based architecture for information-centric vehicular network", *Elsevier Computer Communications*, pp. 133–146, 2017. doi: `10.1016/j.comcom.2017.07.013`

5. F. Modesto and A. Boukerche, "An Analysis of Caching in Information-Centric Vehicular Networks", *2017 IEEE International Conference on Communications (ICC)*, pp. 001–006, 2017. doi: `10.1109/ICC.2017.7997019`

6. F. Modesto and A. Boukerche, "A Novel Service-oriented Architecture for Information-Centric Vehicular Networks", *19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, pp. 136–139, 2016. doi: `10.1145/2988287.2989178`

Acknowledgements

I want to express sincere gratitude to my advisor, Prof. Boukerche, for the continuous support, for his patience, motivation and immense knowledge. I am also thankful to Prof. Boukerche for having allowed me to pursue my studies, accepting me in his lab as well as for his financial assistance, which was critical in enabling my research and allowing me to participate in conferences and publish my work. I also wish to thank, Prof. Soumaya Cherkaoui, Prof. Burak Kantarci, Prof. Iluju Kiringa and Prof. Richard Yu for participating as examiners in the thesis defence committee as well as for their thorough review and insightful suggestions. I want to thank my family for their never-ending support. Adriana, Thiago, Luiz and Luiza, thank you. My dear friends, who are too numerous to name, were also fundamental in allowing me to achieve my dreams and for this I thank them. Most importantly, I would like to thank my spouse, Giovanna, who is my anchor to reality, who has kept me sane throughout my studies and who gave me strength when I had none. Lastly, I would like to thank the following institutions who supported my research: The University of Brasília, the University of Ottawa, CAPES, NSERC DIVA and NSERC TRANSIT research programs.

Contents

List of Tables

List of Figures

Glossary of Terms

BS Base Station

CCH Control Channel

CCN Content-Centric Network

CEE Cache Everything Everywhere

CDN Content Delivery Network

CS Content Store

DONA Data Oriented Network Architecture

DSRC Dedicated Short-Range Communication

Dist Prob Distance Probabilistic

DProb Distance Probabilistic

DTN Delay-tolerant Network

FIB Forwarding Information Table

ICN Information-Centric Network

ICNRG Information-Centric Networking Research Group

IoT Internet-of-Things

IoV Internet-of-Vehicles

ITS Intelligent Transportation Systems

IVC Inter-Vehicular Communication

LCD Leave Copy Down

LDB Load-Based-Deprioritization

LRU Least Recently Used

M2M Mobile-to-Mobile

MCD Move Copy Down

MG-Cache Minimum Gradient Cache

MGRS Military Grid Reference System

MPCS Mobility/Popularity-based Caching Strategy

NBRP Name-Based Routing Protocol

NC Never Cache

NDN Named-Data Network

PIT Pending Interest Table

PKI Public Key Infrastructure

PLC Provider-based Location Correlation

PMF Probability Mass Function

Prob Copy with Probability

PTS Public Transit System

QoS Quality of Service

QoE Quality of Experience

RCOne Random Copy One

RLC Request-based Location Correlation

RSU Road-Side Unit

SaaS Storage-as-a-Service

SCH Service Channel

SEVeN Service-Exchange

SM Service Manager

SUMO Simulation of Urban MObility

TTL Time-To-Live

UG-Cache Utility Gradient Cache

VANET Vehicular Ad-hoc Network

V2C Vehicle-to-Cloud

V2I Vehicle-to-Infrastructure

V2N Vehicle-to-Network

V2P Vehicle-to-Pedestrian

V2V Vehicle-to-Vehicle

V2X Vehicle-to-Everything

VoD Video on Demand

WAVE Wireless Access in Vehicular Networks

WSMP Wave Short-Message Protocol

XaaS Anything as a Service

Chapter 1

Introduction

As vehicular technology develops, networking capability and associated requirements increasingly become limiting factors. Increase in sensor deployment, services provided to passengers and impending vehicle-to-vehicle communication systems are all factors that contribute to increasing reliance on network services. Vehicles are becoming mobile agents with increasingly vast processing power and are expected to be capable of participating in a wide variety of applications. As vehicles become connected service exchange hubs, they will presumably require increased network connectivity systems. Current communication solutions for mobile environments are based on host-to-host architectures such as the TCP-IP protocol stack. However, changes in network structure and energy limitation have long been shown to be limiting factors for such solutions in mobile environments [1]. Many changes have been made in an attempt to solve the issues that have risen over time including the introduction of Mobile IP and proposal of wireless-centric content delivery schemes [2]. However, the core technology behind the Internet was not designed with mobility or content distribution in mind, meaning these solutions are ill-fit and inefficient, increasing the complexity of the Internet architecture [3]. On a broader scope, many recent trends in networking correlated with vehicular networks are also hindered by the limitations of host-to-host. Service-based architectures, distributed systems from Intelligent Transportation Systems (ITS), industrial environments or even those belonging to the Internet of Things (IoT) face limitations in the current host-centric environment. The proliferation of interconnected information exchanging devices and services have requirements that are not met by the current Internet [4], and can benefit from moving towards more localized communication systems [5]. IP does not reflect the optimal format for Internet communication in highly mobile network, and poses a limiting factor in the long run development of Vehicular Networks. Hence, continuing to patch improvements to an inadequate architecture with underlying incompatibilities is not a sensible long-term strategy. Applications belonging to ITS and IoT are prime candidates and will benefit from more flexible environments.

Information-Centric Networking (ICN) is one such architecture poised to replace host-to-host based approaches [6]. ICN is more aligned with the current state of the Internet, one in which communication is increasingly based around information dissemination rather than pair-wise communication [4]. The communication process of ICN is informed by the content objects themselves rather than by the network location where they are expected to be available. By decoupling content from addressing, ICN systems are more flexible and allow for increased mobility and distributed caching support. Among its benefits, ICN allows for services to be decentralized. Decentralization is essential in systems where information is not expected to be obtained from predetermined participants such as Vehicular Ad-hoc Networks (VANETs). As will be discussed in the following chapters, the potential for improvement in caching, mobility support, safety and content discovery are some of the main driving factors favouring the adoption of Information-centric approaches. ICN-based solutions are also significantly more prepared for service exchange, a result of the focus on content and overall increased network awareness. The many improvements of ICN over host-to-host make a good case for its application in vehicular environments. The limited nature of networking resources coupled with the growing demand for services and further motivates the drive for more efficient communication. Because of the high level of mobility, in order to increase performance and supply the ever-increasing demand for content, VANETs will necessarily require more efficient communication architectures than those provided by IP. The mismatch between host-to-host solutions and the unique peculiarities of vehicular networks are driving factors behind a recent research trend suggesting ICN be considered as an alternate approach. The potential of ICN and its better-fit to vehicular networks form the reasoning for this approach being selected for the work that will be presented this Thesis.

1.1 Thesis Statement

Research Problem: Deployment of real-world vehicular communication networks is imminent, yet there are many open challenges yet to be overcome to enable implementation of real-world services. One of the most significant challenges is in enabling vehicles to exchange services, not only sharing information but producing and consuming content generated within the vehicular network. In the face of ever-increasing demand for content, one big challenge in VANETs is how to structure vehicular communication to allow for a wider variety of services while providing minimum Quality of Service (QoS) and prioritizing specific applications such as those safety-related.

Key Idea: In this Thesis, we propose the design of the SEVeN information-centric architecture, its service-oriented components, naming scheme and service prioritization policy, designed to increase communication efficiency and enable service exchange between vehicles.

The implementation of ICN-based solutions for VANETs, including but not limited to network-wide architectures and specialized components, protocols and policies can mitigate the observed issues in host-centric solutions. Moreover, the change in consumer behaviour towards content foments the development of content-centric solutions, a natural fit for ICN. In the context of vehicular networks, constant topology changes significantly hinder host-based routing. A pivot towards content-centric solutions allows for significantly improved content distribution. Solely in terms of caching, the potential for intermediate caching is incomparable, with coordination allowing for decentralized cooperative caching policies not possible in host-to-host scenarios. The implementation of caching policies helps mitigate contention for network access by reducing the distance to content, thus reducing network load. Increased access to content also drives as a motivator for participation in collaborative knowledge exchange networks, fundamental in vehicular safety and information applications. However, individually, contributions are limited by the architecture to which they belong. Towards the goal of providing better mechanisms for content exchange in vehicular networks, in this work, we propose a service-centric ICN architecture, evaluation framework, content access mechanisms and caching policies to study ICN and the many facets in which it is predicted to improve communication in VANETs.

1.2 Objectives

The main objective of this Thesis is to further the research on ICN-VANETs through proposed frameworks and protocols designed to improve on observed limitations or research shortfalls. Specifically, the goal of this body of work is to design a communication architecture that can enable service exchange in vehicular networks. The current stage of research in the field of ICN-VANETs allows for the introduction of novel components at both architecture and component levels. As the design of a complete system architecture very comprehensive. Hence, for the contributions targeting architecture level, our goal is to improve upon existing solutions based on a discussion of likely application scenarios. In regards to contributions made to specific components of communication architectures, we aim to design cooperative caching solutions that are aware popularity-related request patterns to reduce average network load and increase content delivery. In terms of its overall contributions, this Thesis aims to provide a general discussion of information-centric communication in vehicular networks and the need for service exchange. The contributions in this Thesis are presented in the following Section.

1.3 Contributions

The key contributions made by this Thesis are listed following in the order they appear in the document.

In-depth Analysis ICN-VANETs: This contribution consists of an in-depth discussion of the application scope and roles in vehicular networks, followed by a survey of existing works, their contributions and current opportunities in the field. As part of the discussion, we pinpoint the shortcomings in ICN architectures for VANETs as well as denote the limitation in intermediate cache solutions tailored for vehicular environments.

Evaluation of Mass Transit in VANETs: This contribution consists of an in-depth discussion and evaluation of public transit and its potential roles as part of vehicular networks. The applicability of public transit is evaluated by simulation-based analysis of real-world public transit data contextualized to the scope of service provisioning and network coverage. From the coverage evaluation, we pinpoint the significant potential for public transit system integration into vehicular networks.

Service-centric Architecture for ICN-VANETs: This contribution consists of a service-oriented architecture for Information-Centric Networking named SEVeN. Given an increasing application space, the change in nature of how users interact using computer networks, and unique challenges intrinsic to vehicular networks, we propose a novel architecture designed to enable service exchange and service management in highly competitive vehicular ad-hoc networks. The proposed architecture is accompanied by a purpose-defined naming policy and service sub-layer as well as a service prioritization policy named LBD.

Implicit Cache Coordination Policies: This contribution consists of additional discussion on cache coordination, content replication and the introduction of a pair of implicit coordination policies for content insertion based on distance and popularity utility gradient. Said policies were evaluated against existing approaches using computer simulation and were determined to contribute to improvements to content delivery metrics.

Cache Policy for Location-Dependent Requests: This contribution consists of the discussion of provider-based location correlation, the definition of approaches to simulate requests where vehicular communication can display spatial patterns based on either provider or requester positions and the introduction of a caching policy based on social cooperation that improves content delivery with Provider-based Location Correlation (PLC) characteristics. The cache policy introduced improved content delivery in scenarios where displaying PLC. The proposed solution also serves as a baseline for additional works on PLC-aware caching.

1.4 Outline

The remainder of this Thesis is structured as follows. Chapter 2 provides an overview of Information-Centric Networking (ICN) and Vehicular Networks as well as potential applications. Next, Chapter 3 presents the general literature review on the topic of ICN focusing on works with tangible applicability in VANETs. Following, in Chapter 4 we discuss the role of public transit in vehicular networks and argue for increased focus on public transit systems. From the context presented, in Chapter 5 we present the SEVeN service-oriented architecture for ICN-VANETs. Chapter 6 discusses the current state of caching in ICN, applicable approaches to VANETs and limitations that need to be considered by proposed solutions. Following, in Chapter 7 we introduce a couple of implicit cache coordination policies designed to disseminate copies of popular content while avoiding over-replication. Next, Chapter 8 presents a framework to evaluate location dependency as well as cache coordination protocol designed to improve content delivery for applications that display provider-based location-correlation. Finally, chapter 9 presents a summary of this Thesis, its contributions and opportunities for future research.

Chapter 2

An Introduction to Information-Centric & Vehicular Networks

In this Chapter, we discuss the general characteristics of Information-Centric Networks (ICN) and contextualize individual components considered in this Thesis into vehicular networks. The topics presented are based on the overall aspects of ICN as initially proposed by the initial works of Gritter et al. and Jacobson et al. [7, 8, 9]. Hence, the content presented in this Chapter is centred on ICN's core concepts, which are presented in terms of the summation of individual contributions in the field. Discussion of individual contributions is presented following, in Chapter 3. The content of this Chapter is structured as follows. In Section 2.1 we present an introductory overview of Information-Centric Networking, its main principles and components. Following, in 2.2 we discuss the potential use-cases for ICN within vehicular networks. This chapter concludes with closing remarks, presented in Section 2.3.

2.1 Introduction

In this section we present the motivating factors that denote the benefits of ICN as a future Internet technology, list and discuss the core concepts relating to ICN, how ICN is part of a trend in research and contextualize the advances in ICN research in general terms as well as how it can be applied in vehicular technology.

2.1.1 Reasoning Behind ICN

Host-to-host communication was designed several decades ago to interconnect large research institutions. In the 50 years since its creation, multiple mechanisms have been proposed to supplement the deficiencies of IPs original design. This trend saw the introduction of Multicast, Content Delivery Networks (CDN), P2P networks, publish-subscribe models and even mobility-related protocols [10]. Such solutions have extended the lifetime of IP and enabled the widespread adoption of the Internet. However, the underlying limitations of IP and its mismatch with the current Internet model pose significant challenges to its maintenance in upcoming decades [11]. The patch-based approach taken continuously increases the complexity of the Internet while mismatching the demand for content, now more commonly based on multimedia content consumption [3].

The concept of Information-Centric Networking is the result of the limitations of the current Internet model. The existing communication Model, referred to as *"where"* oriented, are generally based on content's physical location such as IP based solutions. In contrast, *"what"* based solutions are centred around content, and it is properties rather than how the data is routed. The priority given to data names is exemplified by a phrase commonly used by works: "Information becomes the first-class citizen". The shift from a *where* model to *what* model was first proposed by Stanford's TRIAD project where it was described as being *"a natural extension of current Internet directory and routing systems"* [7]. Initial publications, dating from 2001, denoted the need for efficient content delivery and based their arguments on the growth of media streaming and exponential popularity growth/demand for content. Formally, initial works tackled content routing and caching, which were described as being limited, heavily dependent on DNS lookups and ultimately not scalable. Gritter et al. coined the term Named Data Networking (NDN) in their work, focused on tackling content routing rather than server or IP address based routing. By contextualizing routing on the desired content themselves, delivery can be significantly improved as data replication in multiple content servers can be abstracted and client requests can be mapped to nearby content servers without the need to contact DNS servers. Initially, ICN was proposed to operate in a symbiotic relationship with existing IP routers. Specialized equipment capable of operating on both IP and the proposed *Name-Based Routing Protocol (NBRP)* would, in theory, parse requests and attach information learned from routing history to improve content delivery, while allowing for existing IP systems to continue operating none the wiser. Later, in 2006, information-centric networking was further popularized by a Google Tech Talk [8].

ICN as a concept has expanded significantly since its inception, and proposed solutions generally moved towards implementing self-contained environments without much consideration to requirements derived from specific applications, though there are works that target this issue specifically.

Information-centric networking is a novel communication paradigm, whose architecture is entirely independent of that proposed by physical addressing solutions. ICN does not rely on an underlying IP system nor does it function as a supplement to improve IP networks. However, ICN can be implemented to operate over an IP network and can naturally incorporate aspects of its functionalities incorporated into its operation. The possibility of ICN being deployed as an overlay solution should come as no surprise, given that IP has matured significantly in its long existence and is the defacto communication paradigm. Many of existing protocols and techniques present in IP networks ranging from flow control to collision avoidance and security can be applied to ICN with ease. Information-centric networks, while designed as a protocol stack functions as an "universal overlay" and can run over any architecture including IP. Similarly, ICN can be used as the underlying service for other protocol stacks. Given the definition of a common namespace, an ICN network could natively cooperate with an IP network and interact with core IP components such as DNS servers.

Vehicular environments are one such context in which the limitations of currently existing solutions are exacerbated, further limiting network service capabilities. It has been demonstrated that IP its mobile counterpart, Mobile IP, is not just ill-suited for the current Internet model but even more-so for mobile environments [12]. The trend of mobile-centric communication, however, is only expected to grow. Cisco has released reports stating that Mobile to Mobile (M2M) traffic is predicted to increase over seven-fold before 2021, while video streaming will account for 82% of all Internet traffic [13, 14]. The mismatch between how users access content and current content retrieval processes limit the quality and efficiency of the services. In theory, ICN is a prime candidate for implementation with the potential to improve content delivery in mobile environments such as vehicular networks. As an example, consider a scenario in which a vehicle is attempting to calculate an optimized route based on live traffic information. In the current Internet context, this is not intrinsically possible without resorting to specialized systems like proprietary mapping applications such as Google Maps [1]. In ICN, however, many vehicles in the region whose information is being requested could serve as a service provider. The client interest is in obtaining traffic information, and they are

[1] maps.google.com

not looking for any given content servers. Assuming all vehicles at the desired location operate under the same assumption of knowledge and trustworthiness, all vehicles in the network can serve as a content provider. The flexibilization of content location is one of many potential examples denoting the benefits of ICN, which can increase content accessibility, reduce the distance to content providers, increase in-network content replication, among many potential improvements. From the discussion of the benefits that ICN can bring to vehicular networks, following we present some of the essential characteristics of information-centric approaches.

One of the main challenges in VANETs is in providing network access to vehicles. The limitations of network resources can be explicitly visualized by considering bandwidth availability. Consider an example of an urban environment where vehicles are equipped with a WAVE network adapter. In this context, it is natural to assume that many vehicles can occupy a compact section of the road, even if the road itself is not limiting vehicle movement. Naturally, vehicles have to share bandwidth with their neighbours, which depending on the network density, can be dozens. Should this be the case, we can demonstrate that the average bandwidth available to a vehicle can become comparatively low. Individual bandwidth availability can be calculated by taking into consideration, road network vehicle density, vehicle dimensions, transmission interference models, network range and bandwidth availability. Following the shadowing propagation model, the power received at a distance χ is given as:

$$P_r(\chi)|_{dB} = 10\log_{10}\left(\frac{P_t G_t G_r \lambda^2}{(4\pi)^2 \chi^\beta L}\right) \tag{2.1}$$

Solving this system for the default configuration of a WAVE adapter: $G_t, G_r, L = 1$, frequency of $f = 5.89GHz$ where $\lambda = \frac{c}{f}$, $c = 3 * 10^8 m/s$ and path-loss exponent β values representing free spaces ($\beta = 2$) as well as shadowed urban environments $\beta = 6$, a receiving threshold of $RX = -89dBm$ results in connectivity ranges varying from 161 to 361 meters depending on transmission power, shown in Table 2.1.1. While the estimates do not consider signal phasing or other attenuation sources, they represent roughly expected connectivity ranges for vehicles equipped with *IEEE 802.11p* WAVE network adapters.

For a series of Sedan vehicle with average sizes of $4.5m$ and spaced every $1m$, a total of 29 vehicles can be included for every road segment. Considering that for this area a single road cross-section is present, in total 58 vehicles per lane are within network interface range. In the scenario considered, all roads in the example region contain two-way streets with two lanes on each direction, totalling 232 potential vehicles in a traffic

TX and RX power values	$\beta = 2$	$\beta = 6$
$P_t = 0.20W$, $P_r = -89dBm$	161m	5.4m
$P_t = 0.28W$, $P_r = -89dBm$	191m	5.7m
$P_t = 1.00W$, $P_r = -89dBm$	361m	7.1m

jam scenario. The increased volume of vehicles in the network would increase the noise level but as transmission power can be increased, so can effective transmission ranges. Ultimately, the number of vehicles within transmission rage will depend on the physical location, relative vehicle speeds and network use. For the sake of completion, consider a scenario with significant interference ($\beta = 6$) in which area of coverages are reduced to $2\% - 3\%$ of their initial range estimates and containing as few as 8 cars. Should the vehicles be equipped with transmitters using the *802.11p* WAVE standard using the standard $10MHz$ bandwidth, they are limited to $18MBps$ bandwidth. This way, each vehicle would have at most $2.25MBps$, and in a worst-case scenario $78kBps$. The low bandwidth available to vehicles very clearly denotes the need for efficient management of network resources, which necessarily requires solutions designed for vehicular networks as is the case with ICN.

2.1.2 Main Components and Terminology

ICN is informed by its focus on content rather than provider location. Information-centric approaches require that certain components exist and adherence to established principles. Note that while ICN is a fresh-start approach, it can be implemented as an overlay atop of currently existing networks, allowing for interoperability with currently deployed systems. However, for the sake of simplicity, for the remainder of this work, we refrain from discussing the incorporation of ICN as an overlay atop of existing architectures. Additionally, we refer to Information-Centric Vehicular Ad-Hoc Networks as ICN-VANETs, omitting the potentially confusing VCCN (Vehicular Content-Centric Networks) acronym in favour of the broader term. Following, we introduce some of the main elements of ICN as a self-contained network architecture.

2.1.2.1 Universal Naming

Naming universality is the core principle of ICN and is defined as the capacity for any device in the network to understand the name of an object. This is because ICN is centred around content identification, meaning that understanding content object names is necessary for network features to function. This does not mean that devices are required to understand the context for the specific object or be capable of speculating on the actual content of the data object. Naming in ICN is transparent and therefore does not rely on association with specific, applications, content providers or their clients. Neither are content object names required to be human-readable. The essential requirement for naming is non-ambiguity, with devices being capable of process named-object requests without naming collision. Naming schemes can be hierarchical or flat. Hierarchical naming is adopted by approaches that use the naming hierarchy as part of content management and routing schemes. While the network is conceptually agnostic to individual names, the components of a prefix can be explored as part of QoS management. Prefix-based service advertisement can also be employed as part of content provisioning. While names are not required to be human readable, individual components of a hierarchical prefix are necessarily distinguishable. Applications can also take advantage of application-based prefixes that are optimal for their context. For example, latitude-longitude prefixes can be employed in applications relating to location services. Flat names, in contrast, offer less flexibility in terms of structuring. However, what flat naming systems lack in flexibility, they offer in trust management, with content names being self-certifiable. In essence, flat content naming schemes can be based on naming encrypted. In this case, only a device with knowledge of the application that provides content objects is capable of producing the names of desired objects belonging to that application. This ensures that content objects are self-certifiable.

Naming is one of the fundamental aspects of ICN. The definition of a unified naming convention is arguably the most critical aspect of any proposed solution. Regardless of the format proposed, the convention defined needs to be so that naming is unambiguous and universally understood. Additionally, all nodes in the network must be capable of naming a request for the desired object as well as being capable of parsing any received requests. Given a solution complies with these requirements it is not bound to additional constraints. The freedom of choice in naming is one of the many benefits of ICN. In VANETs, naming becomes especially relevant, given the connectivity limitations present.

Content names are *opaque* within the network context, meaning that while requests are universally understandable, they are not bound to a specific context or location. Routes do not know the meaning of names and are only capable of assessing their structure. This allows for applications to specify their own naming schemes, enabling them to achieve improved performance while allowing for naming schemes to change and evolve unbounded by network requirements. Application specific naming also allows for improved request batching and granularity control. Content correlation can be employed to improve content delivery. Content providers can batch predicted content along with requested data in order to improve content delivery. Clients can also take advantage of content correlation by making requests with more general prefixes.

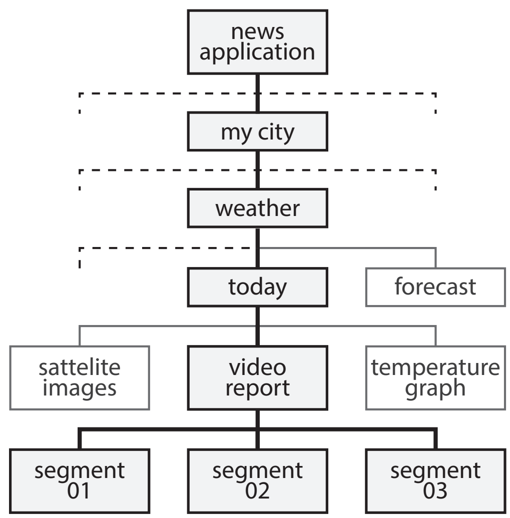

Figure 2.1: Weather Service Hierarchy Naming Example

Consider an example of a multimedia delivery service whose naming scheme is hierarchically structured as displayed in Figure 2.1. In this example, consider a user request for a video report of today's weather by requesting the following prefix: *"[newsapplication.mycity.weather.today.videoreport]"*. The client application can make assumptions about user behaviour and generalize the request to include other useful information by reformatting the request as: *"[newsapplication.mycity.weather.today]"*. In this scenario, by taking into consideration that user behaviour results in further data requests, the client application can preemptively generalize additional content to avoid unnecessary overhead and latency in content provisioning. Conversely, if a system identifies that users only consume the first segment of a video they can restrict requests in order to reduce network bandwidth overhead. In this example, the prefix would be: *"[newsapplication.mycity.weather.today.video.segment01]"*. Partial content consumption is a real-world issue that many service providers especially multimedia service providers have to deal with. One such example is the consumption of video in social networks like Youtube and Facebook. The nature of these services makes it so that users can partially view videos.

Notably, video advertising is not relevant for users and is generally skipped, making partial video caching a relevant feature in this scenario. By taking advantage of namespaces in ICN, content can be structured so that delivery and caching are optimized.

As names are not required to be meaningful in a network context, naming opacity allows applications to structure their namespace in order to include information that cannot be understood by outside viewers. Applications can also choose to format their requests in non-readable or binary formats in order to improve namespace efficiency. This allows for efficient lookup techniques such as bloom-filters to be applied during cache lookups. Other formats such as binary naming can also be employed to improve lookup for popular items. Granularity is yet another aspect of naming that requires efficient design. Unrestricted naming can become a liability and cause adverse effect to packet header lengths and lookup times. This should come as no surprise, as prefix sizes can increase overhead if unbounded and hierarchical tree prefix traversal is proportional to the length and depth of the requested item. For example, in the IEEE WAVE's 160.3-2010 standard, content provider service context header extension fields are limited to 31 octets and up to 32 service info segments can be added to a single Service Advertisement, which cannot exceed 255 octets [15].

2.1.2.2 Content Discovery and Delivery

The pivotal change towards universally understood named objects results in a departure from the philosophy adopted in host-to-host systems. Different from Host-to-host communication, requests are not directly targeted towards specific host locations. Most service discovery strategies available perform poorly in Vehicular Network environments [16]. The underlying concept behind content discovery and routing is reliant on establishing pathways toward content sources. The decoupling of content from its physical address entails significant changes to how requests are fulfilled. As content is not tied to a specific location, information-centric solutions rely on different content-discovery mechanisms. The approaches for content-discovery are divided into client-driven or provider-driven approaches with routing being designed around these two diverging methodologies. In the provider-driven approach, content and service availability is broadcast by producers. Peers connected to these content servers store DNS-like routing tables listing which content providers are accessible at specific network interfaces and what content can be accessed through each interface. This role can either be performed by all members of the network or be assigned to specialized network entities. Content discovery is tied to individual requests and functions based on requests or registration for available services. The knowledge regarding available services allows for subscription-based content obtention. In client-driven approaches, content discovery is initiated by clients, based on interest propagation. Content availability is not necessarily known at the time of a request for content and relies on collaboration with neighbours. Communication is divided into a minimum of two types of messages, with requests for content being broadcast as "interest" messages and replies with the requested content usually being labelled as "data" messages. Interest request propagation can be made by broadcasts or unicasts targeting specific nodes. Regardless of which propagation mode is elected, a certain level of flooding is expected as part of interest propagation. Routing in client-driven approaches usually follows a reverse-path from the content source to its original client. The trade-off between provider-driven and client-driven approaches relates to the overhead costs associated either with maintaining content routing tables or those caused by flooding. In both instances, it is relevant that the communication model be as loosely coupled as possible, in order to adapt to the highly dynamic network structure of vehicular networks.

2.1.2.3 Caching

Caching is one of the key components of ICN whose pervasiveness differs significantly from traditional host-to-host. While technically not a requirement in ICN, caching is one of the areas with the highest potential for improvement in ICN. Naming universality allows for significant expansion of caching functionality, mainly in terms of improved in-network caching. All nodes of an ICN network can operate as in-network caches, allowing for deployment of more efficient caching policies. This potential for use drives the expectations associated with cache policy design and deployment. Collaboration between devices, predictive caching, and social caching are amongst the many potential techniques that can be employed by ICN networks. The challenges associated with vehicular networks make caching an even more critical aspect of ICN, as the high level of mobility of vehicles increases the need for retransmissions due to network topology changes. By efficiently storing intermediate copies closer to the position of the requesting vehicle at the moment of its request, there is an increased likelihood that at least a portion of the route towards the content provider will remain the same. In this case, a request will not need to be propagated the full path towards their original provider, being fulfilled by an intermediate node that previously participated in the process. In essence, there are many caching approaches that can be incorporated in ICN-VANETs to improve content delivery. Proactive caching solutions, for example, can be employed to place content at strategic locations within the network based on request pattern prediction or other heuristics.

2.2 ICN for Vehicular Systems: Use Case Scenarios

The Internet is, for the greater part, a network driven by content consumption. The vast number of applications and continuously evolving use-cases in networking have lead to the development of an agnostic network structure. Vehicular networks, on the contrary, can be clearly divided into the general and a second context, one based on in-network service exchange. Not only can vehicles provide services directly to their peers but some applications, especially those relating to safety and mobility are beneficial to the network as whole and candidates for prioritization. Nevertheless, vehicles are predicted to implement an increasing volume of applications in varied contexts. As the potential application-space increases, it becomes relevant to understand the potential ramifications of each use-case. While it would be futile to list uniquely specific applications, we can discuss them in terms of their general characteristics. Depending on the application context, vehicles may operate as providers, clients or as both. Application contexts also vary depending on the actors involved, with all potential contexts being encompassed under Vehicle-to-Everything or V2X for short. Interactions with other vehicles, with infrastructure or services, dedicated cloud services and other actors including pedestrians are labelled under different groupings depending context and involved actors and are referred to Vehicle-to-Infrastructure (V2I), Vehicle-to-Cloud (V2I), Vehicle-to-Network (V2N), Vehicle-to-Vehicle (V2I) and Vehicle-to-Pedestrian (V2P) respectively.

Interactions classified as Vehicle-to-Infrastructure (V2I) represent those where vehicles are communicating with infrastructure-connected devices. Grewe et al. suggest distinguishing interactions with Infrastructure Edge services from The Internet by classifying Communication with Internet-connected services as Vehicle-to-Cloud (V2C) [17]. However, for the sake of simplicity, and as many edge services are reliant on asynchronous connection with Infrastructure, in this Thesis, we refer to all infrastructure communication as belonging to V2I. We present an extensive series of potential applications within Vehicular Networks that can be made available as part of ICN systems in Tables 2.2 and 2.2. Applications are grouped based on their context and category represent the applications considered by the literature, many of which are listed in [18]. For each application, we present a brief description that is elaborated following as well as our classification for this application. Given the large variety of services that can be provided by vehicles to their neighbours, all applications classified as services are labelled under the collective *XaaS*, shorthand for *Anything as a Service*. Note that there are equally-named applications listed under multiple contexts. This is due to how applications are structured, with applications behaving under different premises and in different formats within these contexts.

Application	Classification
Safety	
Accident Report & Warning	Emergency Events
Road Hazard Report & Warning	Emergency Events
Evacuation Coordination	Emergency Events
Emergency Network Coordination	Emergency Services
Emergency Public Service Advertisement	Emergency Services
Knowledge Sharing	
Parking-lot Status & Assistance	Public Services
Electronic Toll/Fare Collection	Monetization Services
Sensor Information Knowledge Sharing	Information Services
Traffic Information Advertisement	Information Services
Route Recommendation	Information Services
Traffic Light Optimization	Information Services
Vehicle-based Support Network	XaaS
Side-channel Information Validation	XaaS
Infotainment	
Video on Demand (VoD)	Multimedia Services
Internet Access	Multimedia Services
Context-based Advertising	Monetization Services
Parked Vehicle Clusters	XaaS
Hardware as a Service	XaaS

Table 2.1: Vehicle-to-Infrastructure (V2I) Applications

A list of applications relating to interactions between vehicles and infrastructure services (V2I) are presented in Table 2.2. Most use-cases in V2I are predicated on infrastructure-based service delivery. The degree of interconnectivity of infrastructure edge devices allows them to provide coordinated access to the Internet as well as to specialized services unique to the vehicular environment. Network-wide infrastructure service access is especially important when considering safety-related applications. Given the mobile nature of vehicles, infrastructure-based crisis and emergency events, infrastructure-based management is beneficial as it allows for simultaneous and continuous vehicle coordination. The distributed spatial awareness of infrastructure services also allows for the management of locally relevant events. In some cases, notably, understanding of network and real-world conditions can be a result of vehicle-generated knowledge. Collaboration between vehicles and infrastructure allow for further improvements in infrastructure-based service delivery.

V2I applications can benefit from ICN integration given the potential benefits ICN offers to content and service-oriented architectures. In the current Internet model, devices have to access a specific set of servers to request content or to become subscribed to services. In contrast, in information-centric environments, content is not tied to any specific server and, therefore, service providers can be distributed, allowing for infrastructure-edge services to operate in a more decentralized fashion. This decentralization means that infrastructure service providers can be placed closer to the locations from where they are accessed. Independent service operation by infrastructure-edge devices decrease the distance from clients to content and increase fault-tolerance resistance. Moreover, increase intermediate-caching capabilities further increase content availability and decrease the distance to content. Both metrics are fundamental, as competition for network services is expected to increase. Service classification can also be taken advantage of, with named data being explored for the implementation of classification-based Quality of Service (QoS) policies. Safety applications are a clear example in which migration to information-centric solutions is beneficial for vehicular networks. Such is the case with DSRC, whose main goal is enabling safety-related communication systems. In the face of increased service availability, safety applications can benefit quality assurance in ICN made possible by explicit application-based policy-making. Applications from different service categories can also be integrated and benefit from ICN-based architectures. Knowledge-sharing services can take advantage of direct reports from vehicles, decoupling knowledge of road network and other sensor-based knowledge. Similarly, vehicles can be contacted to validate information provided to the infrastructure. For example, accident report status updates, photographic evidence can be requested by infrastructure services directly to vehicles in regions of interest. Public transit is a crucial player in this category of application, with additional discussion presented in Chapter 4. Nevertheless, applications outside the scope of cooperative and common interests are also possible as part of V2I scenarios. Infotainment services can take advantage of dedicated network infrastructure for its connectivity to the Internet. The potentially closer distance to infrastructure edge devices such as RSUs and the services that can be provided by these devices can improve service quality. Ranging from infrastructure-side edge caching to an overall increase in connectivity, infotainment services benefit from interaction with infrastructure. Given the higher bandwidth requirements of many multimedia applications, increased network access is a necessity to enable some applications to function. On-demand Video is an excellent example in which the application requirements mandate reliable network access. Other applications with similarly high demand for bandwidth or with strict latency restrictions can also be enabled by vehicle-to-infrastructure connectivity.

Vehicle-to-Vehicle (V2V)

Application	Classification
Safety	
Emergency Brake Warning	Emergency Events
Obstacle Warning	Emergency Events
Pre-collision Warning	Emergency Events
Road Hazard Warning	Emergency Events
Wrong-way Warning	Emergency Events
Lane Change/Departure Warning	Mobility Services
Left/Right Turn Warning	Mobility Services
Overtaking Warning	Mobility Services
Priority Vehicle Approach Warning	Mobility Services
V2V Mobility Collaboration	Mobility Services
Knowledge Sharing	
Traffic Information Sharing	Information Services
Weather & Sensor Sharing	Information Services
Point-of-Interest Information Sharing	Information Services
"See-Through" Live-camera feeds	Information Services
Anything as a Service	XaaS
Infotainment	
Anything as a Service	XaaS
Video and Audio Conference	Multimedia Services
Video on Demand (VoD)	Multimedia Services
Connected Video-Games	Multimedia Services

Table 2.2: Vehicle-to-Vehicle (V2V) Applications

While V2V systems benefit from ICN environments, the true potential of ICN is observed in the vehicle-to-vehicle application space. V2V communication is heavily reliant on the efficiency of routing. Because all vehicles can potentially take the role of content producer as well as perform intermediate caching, there is a natural increase in content availability. This, in turn, allows an increase in service quality and introduction of a series of novel applications. Naturally, vehicular collaboration and knowledge sharing has been one of the driving factors behind the implementation of vehicular networks. The quality of safety applications is heavily dependent on latency. Hence, by allowing nodes a more ample role in the content dissemination process, safety services increase in efficiency, potentially saving lives. Contextual awareness of the application class can be taken advantage off, enabling vehicles to prioritize specific subsets of applications to ensure that critical services are provided regardless of network load. Such prioritization schemes can be performed without specific knowledge of individual services specifically because naming systems in ICN can offer insight into applications. Similarly, the introduction of additional services to an existing V2V network will require little change, as they should be self-descriptive. As the volume of sensors in vehicles increase, the number of safety applications they provide an increase, with additional warning systems being introduced. Regulation of safety systems and their behavior will need to be put in place and updated as new applications become available. Nevertheless, vehicles that do not understand the context of novel safety applications will still be capable of forwarding related knowledge without needing to understand the detailed context of the applications in question. This logic applies to all categories of services in ICN. A list of frequently cited applications classified according to their role is presented in Table 2.2. As discussed, safety-related services are the driving factor behind VANETs. Consequently, the safety-related application-space is well-established but expected to grow. Currently predicted safety applications include emergency event warnings and other critical awareness applications. Note that, we also include mobility collaboration as a safety application. This was done as a means to describe contexts where vehicles must coordinate their mobility in order to ensure safety. Lane merging is one example of mobility coordination that falls under the safety grouping. Other, non-safety-critical coordination applications fall under knowledge sharing schemes. Knowledge sharing systems describe applications that are not safety-critical, but that benefit all members of the network. Unlike safety applications, knowledge sharing systems are unlikely to be regulated or mandated by the government. Hence, knowledge sharing systems are reliant on proactive cooperation to function. Naturally, VANETs can implement systems to incentivize knowledge sharing, and the outcome of such systems is beneficial for participants, which further incentivizes participation. Decoupling knowledge systems from proprietary sources, as exemplified earlier in the Chapter, is yet another reason to develop such applications. The ramifications of knowledge sharing go beyond information and include collaboration in routing, caching and all hardware and logical resources. XaaS can be explored as a tool to enable collaboration schemes, and the very mechanisms can be interpreted as being part of the collaboration process. For example, forwarding of content can be observed as an altruistic gesture that benefits the network. This way, all aspects of vehicular technology can be incorporated under the knowledge sharing umbrella. XaaS can also be explored with the intent of benefiting oneself. In this case, service provisioning is characterized as part of infotainment services and is provided either for personal financial gain or as a currency necessary to participate in a service exchange system. Other infotainment applications are more clear-cut, relating to the personal interests of passengers. This includes entertainment applications and non-critical information systems.

In-vehicle services generally fall under the umbrella of infotainment and can be provided either by V2I or V2V systems. Similarly, other applications can take advantages of both infrastructure and ad-hoc communication models. Collaboration with previously listed entities, such as passengers or cloud services, can further improve the quality of the applications available within a VANET. Nevertheless, V2I and V2V are the two primary contexts and, for most scenarios, can be used to classify specific contexts within V2X systems.

2.3 Concluding Remarks

In this Chapter, we introduced ICN, provided a brief contextualization of its inception, its benefits relating to vehicular networks. Following we discussed the core concepts of ICN, introducing and classifying the potential approaches taken to tackle individual aspects of ICN. Then, we proceeded to present and discuss a series of scenarios and associated use-cases specific to vehicular networks. This discussion and the associated examples serve as a basis for the discussion that will be made following in Chapter 3, where we analyze the current state of the art and the contributions made by the existing literature. Similarly, the fundamental concepts presented in this Chapter are a point of reference for the discussion on public transit presented in Chapter 4 as well as for the solutions proposed in Chapters 5, 7 and 8.

Chapter 3

Literature Review

The content of this Chapter is structured as follows. We introduce the existing literature in Section 3.1. Section 3.2 follows, discussing the literature review in terms of different components of ICN. This chapter concludes with closing remarks, presented in Section 3.3.

3.1 Introduction and Context

First theorized in 2001, ICN is a novel approach to network communication, introduced as part of the TRIAD Project [7]. Noteworthy publications were first published starting in 2009 [9, 19]. Significant attention in the field has been paid to mobile ad-hoc networks, as ICN can take adventure the particularities of such environments. Vehicular networks were one of the theorized applications. However, research specific to vehicular networks is recent and still in an early state. In this Section, we discuss the current state of ICN through the outstanding works in the field, focusing on contributions with applicability in vehicular networks.

3.1.1 Current Advances in ICN

ICN, while still in its infancy, has gained significant traction in the past five years. Starting in 2008 several research groups began working in the area, and a select number have expanded into very productive research groups. The various resulting architectures, now referred under the umbrella term Information-Centric Networking, follow a basic set of principles as detailed in Chapter 2. The philosophies of these approaches vary significantly, with individual solutions focusing on specific functionalities or challenges. The inception of these approaches is also a defining factor, with most notable initiatives being either developed and managed by universities or by governmental agencies. Research projects we will present following are either managed by academic institutions or partnership with governmental agencies, such as the European Research & Innovation Commission FP7's supervision. The most notable projects relating to ICN are as follows:

DONA: The first major ICN architecture introduced, Data-Oriented Network Architecture (DONA) is the initial result of the 7th Framework Program (FP7) [20]. The FP7 was the European Union's Research and Innovation Funding Programme for 2007-2013. Many of its projects are ongoing, and the current iteration of the funding program is called Horizon 2020. DONA proposes a change in addressing scheme, separating content identifiers from locators. Names in DONA are self-certifiable, and Service and data discovery in DONA are dependent on *Rendezvous Network* (RENE) systems, which function similarly to DNSs. However, there is no rigid structure in the RENE, a consequence of the change to named-data. Content provisioning follows a Publish/Subscribe model with Registration messages being used to maintain registration lists. Many research projects followed DONA's approach [21]. 4WARD, is one such approach, is designed around a "Strata Concept", with a vertical stratum correlating information from different horizontal strata taken into account network resources and characteristics of underlying network architectures [22]. Technically, it is designed to allow for enforcement of multi-layer traffic engineering policies in multi-homed devices. The stacking of horizontal strata is designed to avoid the strict layering of the OSI model, taking advantage of the unique characteristics of each horizontal strata. Scalable & Adaptive Internet Solutions (SAIL) and subsequently Network of Information (NetInf) followed 4WARD. NetInfo, the Work Proposal 6 of the 4WARD project, is tasked with designing the network layer for the architecture.

PIRSP/PURSUIT: The Publish-Subscribe Internet Routing Paradigm (PSIRP), was one of the first projects supported by FP7 and has been superceded by PUblish SUbscribe Internet Technology (PURSUIT) and more recently Line Speed Publish/Subscribe Inter-Networking (LIPSIN). PIRSP and its following projects propose the implementation of publish/subscribe systems for ICN content. The proposed architecture has a functional separation of functionality between content access and content delivery, both interacting with a component dedicated to topology management to ensure content delivery. Much like in DONA, a *Rendezvous*

Network is tasked with name resolution and routing. Publish and Subscribe messages are used to maintain registration lists, with a topology manager being tasked with ensuring that content is routed towards subscribed nodes.

Mobility First: The MobilityFirst architecture is focused on the challenges associated with mobile Internet scenarios. However, Vehicular networks are among the scenarios in which the research group focuses. MobilityFirst-NP is the successor to Mobility First. This second phase of research departs from its predecessor in terms of its goal. MobilityFirst-NP is designed to operate in service-centric environments focusing on use-cases such as the Internet-of-Things. The research of this project is mainly focused on mobile and cellular systems including 4G and 5G.

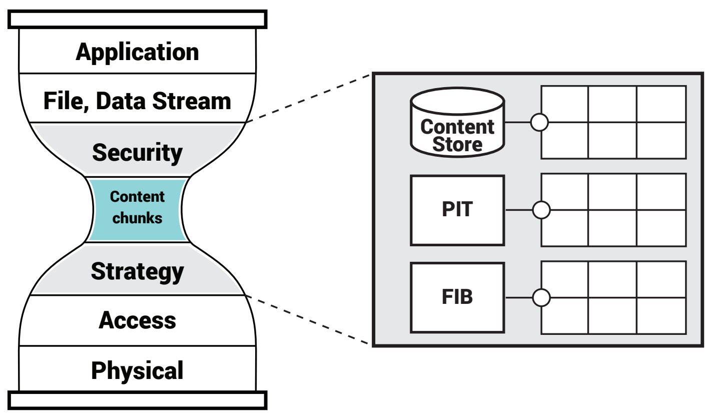

Figure 3.1: NDN Hourglass Model

CCN/NDN: Content-Centric Networking (CCN) is the prevailing ICN architecture with significant support from the research community [8]. CCN was described during a Google Tech Talk and later formalized by works from Jacobson et al. and Xerox PARC. The Named Data Networks (NDN) approach, as proposed by Jacobson et al [9] was the first full-fledged architecture proposed based on CCN. The CCNx protocol and simulation architecture proposed by PARC derived from Jacobson's initial work and is the bases of current Named Data Networks [19]. The pervasiveness and the entangled manner in which research on CCN developed frequently leads to confusion among authors, with ICN, CCN and NDN being used interchangeably. Regardless, NDN is considered as a clean-slate architecture, replacing host-to-host based solutions. Notably, NDN systems can be implemented as an overlay solution to allow for interoperability with existing systems. NDN proposes a hierarchical Namespace, similar to that of Uniform Resource Identifiers (URIs) and cryptography is achieved using a Public Key Infrastructure (PKI). Named requests are made hierarchically, and routing leverages aggregation through topology-independent naming scheme to avoid duplicate requests. NDN retains the thin-waist moded used by the IP model and introduces a strategy layer, as presented in Figure 3.1. Its proponents argue that the simplicity of the core layers allows for service flexibility. By removing the restriction that packets can only name communication endpoints, NDN is designed to reshape the architectural framework of the Internet. This fundamental difference from IP allows for consumer-driven data delivery with only the minimal required functionality. NDN also introduces a novel packet forwarding engine based on client-driven requests. The NDN structure is divided into three components, the Forwarding Information Base (FIB), the Content Store (CS) and the Pending Interest Table (PIT). The FIB is responsible for storing known potential sources for content and stores a list of outgoing faces associated with different content objects. The CS functions as a cache and stores content objects both for incoming requests as well as during content routing. The PIT contains a list of currently pending requests and associate requesters, both local and remote. Upon content being obtained, its compared against the PIT to solve pending requests. From the vast body of literature, many network simulation tools have been developed to study CCN and its derivative works NDN has a custom simulator (ndnSim) based on the NS-3 simulator. The CCNx distribution has evolved significantly since its introduction, and its current iteration has developed into a modularized simulation framework. Conceptually, modules can be developed independently and integrated as additional features to the CCNx core. Alongside its main core (Libccnx-common), the current architecture includes several other modules. CCNx Distillery 2.0 is a software distribution system that is responsible for combining different modules and assembling the CCNx Software Suite. The NDN Forwarding Daemon (NDN-DF or NFD) follows the basic NDN Concept but diverges from the CCNx approach in favour of implementing its core components and application libraries to increase overall performance. Cicn, a fork from the 1.0 implementation of the CCNx 1.0 implementation was developed as part of the research effort made by the IRTF ICN Research Group. Lastly, NDN-NP has carried on the concepts proposed by CCN but abandoned the CCNx Prototype developed by PARC, focusing on its NDN-based prototype.

3.1.2 Information-Centric Networking Surveys

Over the years, several surveys focused on ICN have discussed current research on the topic of ICN with varying scope and depth. Most have focused on specific topics and associated challenges while others have made more broad discussions of state of the art. In [23], the authors present one of the very first surveys in the field. In their work, they introduce the common basic concepts of ICN and compares a body of work in terms of ICN's core elements, introduced in Chapter 2. Similarly, [4] compares different ICN architectures based on the core components as well as introducing a discussion on security. [5] discusses present challenges in ad-hoc networks, promoting the adoption of ICN for such environments. [6] then focuses its discussion of ICN on wireless networks and the associated challenges. The authors evaluate the applicability of the CCN approach in terms of node mobility and resource constraints. CCN is the main focus of yet another work, with [21], performing an in-depth discussion of solutions derived from the CCN approach. This work explores the suitability of a subset of approaches to VANETs. [24] performs a short survey of challenges relating to CCN in the context of vehicular networks, focusing on VANET centric paradigms. [25] discusses mobility support in infrastructure-based NDN, classifying existing solutions based on their design characteristics. The work in [26] focuses on the challenges associated with mobility and vehicles, covering Connected Vehicles as a concept within ICN. It provides analysis that supports the theory behind adopting ICN in vehicular networks. In [27], the authors discuss Interest Forwarding mechanisms in client-driven Vehicular ICN. [28] performs an analysis of caching and forwarding in ICN. The authors propose a taxonomy model to classify caching policies, discussing the advantages and challenges of different approaches. Both [29] and [30] survey the existing literature in ICN caching policies, focusing on evaluating and discussing caching redundancy and performance metrics respectively. [31] presents a survey on the security and privacy aspects of ICN, one of the main open challenges of the field. They discuss the existing body of work in the area in terms of their underlying principles and drawbacks. [32] also tackles security in ICN, surveying attack vectors unique to ICN. Recently, given the natural expansion of the field, a couple of relevant long-form surveys have been produced. [33] discusses the general topics of ICN, the evolution of the research field, introduces a stage-based classification for academic research on the field and discusses individual characteristics of ICN. In this work, the authors make a note of the applicability of individual concepts and approaches of ICN, denoting those either applicable or targeted towards vehicular networks. In [17], unlike most works, where general applicability of ICN is discussed, the authors dissect the ICN architecture in vehicular networks. From this, they denote several application contexts and open research questions specific to the context of automobiles.

3.1.3 Information-Centric Vehicular Networks

Information-centric is considered by a more fitting approach for vehicular networks, given the improvements that it offers in comparison to host-to-host solutions relating to challenges in mobile and ad-hoc environments. The Internet Research Task-Force ICN Research Group (IRTF ICNRG) has explicitly presented VANETs as one of the baseline scenarios for its application in an RFC [18]. VANETs are defined by the peculiarities of their environment, namely: (1) unlimited energy resources, (2) ample processing power, (3) large storage capacity, (4) high level of mobility, (5) limited and intermittent connectivity and, (6) specialized service requirements. VANETs are not limited by their energy requirements as vehicles produce energy while in use. Therefore, solutions are not necessarily required to consider energy consumption as one of the most significant limitations. This contrasts with MANET, topology changes as a result of power failure is an ongoing threat. It also means that more powerful communication devices can be deployed in vehicles. The inclusion of more powerful network interfaces is also relevant when dealing with other issues in VANETs, such as mobility and intermittent connectivity. For any level of QoS to be implemented, minimal service requirements have to be met, which can be a significant challenge in VANETs [34]. Vehicles in urban scenarios have to deal with very high levels of interference result of signals bouncing off buildings and the volume of vehicles that can be connected in small chunks of road. Highway scenarios pose other threats, with vehicle sparsity and minimal connectivity to infrastructure limiting access to content. Both examples denote the benefits of the deployment of more capable devices, capable of inter-connecting multiple networks via separate interfaces, and better power control. Ample processing power allows for VANET applications like clustering and other functionalities to be made available. Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS) as well as Collaboration as a Service (CaaS) can take advantage of ample processing and storage capabilities. Some works target vehicles in parking lots as a processing and storage source. VANET potential for knowledge and service exchange is another important factor that favours ICN.

The development of information-centric solutions specifically tailored for vehicular networks has trailed overall research on ICN. While many prominent research projects have highlighted VANETs as a key application scenario, most have forgone in-depth research on specific topics in favour of establishing baselines and fine-tuning architectures. Most projects belonging to the initial stage of research on ICN, as classified by Liu et al [33], focused on overall system architectures. From this research basis, more recently, significant focus has been given to vehicular networks as applications of existing proposals for ICN. One unique challenge associated with vehicular networks relates to the longevity of the technologies deployed to vehicle. Unlike enterprise and consumer-grade products, which can be independently upgraded, vehicles are assembled as complete packages. This means that most networking elements deployed are, not only unlikely to be replaced during the vehicle's operational life-cycle but also that because vehicles can be kept in operation for considerably long periods, whatever solution is deployed needs to have long-term compatibility with newer systems. As vehicles have operational durabilities span from 10 to 15 years, solutions need to compensate for over a decade of technical innovation [35, 36]. Different vehicle manufacturers tend to employ different sensors and incompatible brand specific middleware tailored to their vehicles, meaning VANET connectivity can be limited, leading to network fragmentation.

3.1.4 Communication Frameworks

The implementation of ICN in VANETs is the subject of significant discussion. The academic community, industry and government have been entangled in a discussion relating the viability of different approaches. The technologies most likely to be implemented are DSRC WAVE and 5G.

3.1.4.1 DSRC WAVE

IEEE's 1609 workgroup, initially referred to as the Dedicated Short Range Communication (DSRC) workgroup is responsible for the Wireless Access in Vehicular Environments (WAVE) Networking Services Protocol family, one of the leading contenders for adoption in ITS systems. WAVE is designed for ITS systems and is considered as the baseline for CCN based solution available in the literature. Its maturity is a result of significant effort in ITS systems, and its flexibility and potential independence from IP architectures allow for its adoption in ICN environments. As described by the IEEE 1609.3-2010 standard [15]: "WAVE provides a communication protocol stack optimized for vehicular en-

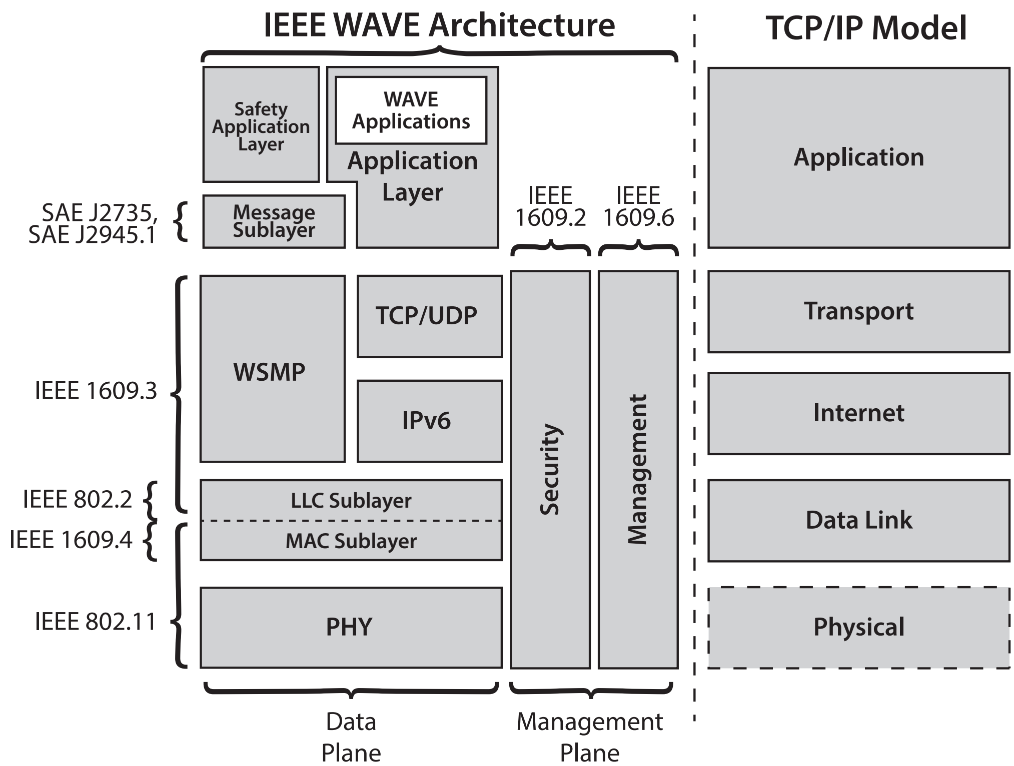

Figure 3.2: IEEE WAVE Architecture (IEEE 1609 Family)

vironment". Figure 3.2 presents the architecture model for IEEE 1609. The architecture is designed to operate alongside TCP/IPv6 and implements its protocol named *Wave Short Message Protocol* (WSMP). WAVE is designed to operate as a producer-driven model, where providers announce WAVE Service Advertisements (WSA) indicating data availability and users, who initiate service requests for broadcasted services. WAVE and the WSMP protocol are a base framework for many ICN frameworks. Hence, ICN can operate in its receiver-driven model atop of a WAVE architecture. In fact, given the flexibility of both platforms, many of the services proposed for WAVE can be integrated into ICN. ICN is highly interoperable with WAVE and can benefit significantly from integrated deployment. One example is in the IEEE 1609.2 standard, which describes the security services as part of the management plane. Annex C of the document lists specifications for applications and can be used as a baseline for an integrated security solution in ICN. The resulting IEEE 802.11p standard (Amendment 6: Wireless Access in Vehicular Environments - WAVE) is designed to operate in the 5.9Ghz bands as a short-range communication system.

3.1.4.2 5G

5G is the latest generation of cellular mobile technology, predicted to be introduced starting 2020 [37]. Unlike DSRC WAVE, whose decade-long development has been embroiled with public rule-making, 5G is a non-exclusive platform that expands on the communication standards used for current mobile systems. Connection bandwidth is one of the main driving factors behind the interest in adopting 5G, with Samsung and Qualcomm's initial trials boasting data throughput over 1.7Gbps [38, 39]. 5G C-V2X connectivity unification, as proposed by Qualcomm, is possible with the evolution of 5G and the expected design for Release 16 (R-16) Phase III version of the 5G Standard. Qualcomm describes this process being implemented based on careful spectrum planning and integration of the established V2X 802.11p, allowing for complementary integrated operation of multiple systems [40]. They suggest that 5G be used to increase bandwidth and application diversity while maintaining support for the existing platform and support for IEEE WAVE and ETSI-ITS. The significant support of the automotive industry, which moved towards integrating 4G systems, is critical and likely to result in 5G is a unifying platform for future vehicles. A study performed by the 5G Automotive Association conducted a performance evaluation of IEEE 802.11p and LTE-V2X (PC5), denoting superior performance of the 5G system due to a combination of increased performance and connectivity range between devices [41]. While the study was performed by an industry association with a bias towards the technology they are developing, the results presented in this work, reinforce the need for integration of both 5G and WAVE standards as expected for C-V2X R-16 design.

3.1.4.3 ETSI

In Europe, the European Telecommunication Standards Institute (ETSI) is responsible for ITS service standardization. Currently being fine-tuned, Release 1 Standard defines the basic principles of operation for VANET communication in ITS environments. The basis of the ITS standard is defined by ETSI TS 102 723-8, which specifies the ITS communication architecture, and is the basis for all other standards [42]. Figure 3.3 displays the predicted applications schedule to integrate the ITS services and includes the proposed ITS-G5 vehicular services. VANET Wireless communication systems for ITS, road transport and telematics are proposed to use frequency bands $5GHz$ and

Figure 3.3: ETSI's Predicted ITS Applications

$63GHz$. ITS architectures are implemented on top of DSRC and CALM and include standards similar to those proposed by IEEE 1609 [42]. For example, the European Standard EN 300 674 - Road Transport and Traffic Telematics (RTTT) is designed to work in $5.8GHz$ ISM bands and is to be used for Electronic Fee Collection, similar to IEEE 1609.11 [43]. Interaction between vehicles and roadside systems is designed to use CALM and operate over the $63GHz$ band.

3.2 Related Work

3.2.1 Naming

Because content names are a fundamental aspect of ICN, there is significant research on the topic. Most works focus on designing solutions tailored to an existing architecture and adhere to the NDN architecture. Contributions range from viability evaluation of existing ICN models to proposal of novel naming and sorting structures specific to vehicular network. In [44], the authors investigate data naming in single-hop V2V and

V2I networks based on the CCN framework. They propose a prefix system structured as follows: */traffic/geolocation/timestamp/datatype*. The proposed model is suggested to take advantage of the spatio-temporal scope of traffic information, and the leftmost component lists the application names. The authors discuss the potential for a typical case where both location and time information are relevant in application contexts with independent granularity control. [45] proposes the use of *Octrees* as data representation structure for 3D environment representation as part of their CarSpeak communication model. Octree encoding is proposed to contain tuples of 8 elements representing the occupancy of the lower level layers of a cube. Cube sections with information are expanded into smaller cubes hierarchically while unoccupied or unknown layers are not expanded. The proposed Content-Centric MAC is designed around "regions" rather than nodes, and the efficiency of the recursive structure used is the result of the compact representation nature, simple encoding technique and efficient bandwidth allocation result of the spacial structure proposed. In [46], the authors propose a bloom-filter based content discovery scheme. The proposed routing scheme is designed to handle the large-sized VANETs dominated by non-sharable data services. The proposed prefix system is structured as follows: */Category/Service_name/Additional_info/*. Categories represent one of three popularity based classification identifiers. [47] introduces a naming Scheme based on Geographical Location. The proposed name-data model efficiently maps bi-dimensional geographic areas into a uni-dimensional naming scheme using a pairing function to map n-dimensional sequences of digits into an ordered hierarchical name that is understandable and easily processed by other nodes. The approach improves in-network caching by reducing prefix sizes. The authors also claim that it allows for directional interest forwarding, as the locations listed can be used as part of a geographical routing scheme. [48] evaluates the applicability of NDN to the vehicular environment. The authors implement a prototype V-NDN network and perform real-world experiments on text and multimedia content. From their experiments, they denote the benefits of named data in allowing for forwarding nodes to handle data based on application needs and increase cooperation by intermediate nodes. [49] introduces a mixed naming scheme based on hierarchical naming that uses a *Compact Trie* (CT). The proposed naming scheme includes both a hierarchical portion and a hash/flat portion. The initial portion of the prefix utilizes a hierarchical structure to identify the application and general properties of the object such as content provider, type and other relevant attributes. Following, the authors introduce a scheme that includes a hash that uniquely identifies content within their respective applications. This approach was taken to improve application-based routing and allowed for the implementation of the CT, which can perform faster lookups than the table model initially proposed for NDN. Similarly, in [50], the authors introduce a mixed naming scheme that capitalizes on the strengths of flat and hierarchical approaches. The model also considers the presence of a hierarchical component followed by flat content identifiers and auxiliary labels. The Hierarchical Routing Prefix (HRP) is used for routing while the Flat Content Identifier (FCI) is used to identify content chunks. Additional attributes are used for data encoding and other element descriptions such as bitrate requirements and priority level. [51] introduce a hierarchical naming scheme specialized for inter-vehicular communication (IVC) based on NDN. The authors argue that the lack of geographical positioning in NDN's naming scheme is a limiting factor for routing. The proposed naming model uses device-based and location-based naming schemes. Inclusion of device name was labelled as relevant to provide identifiers, necessary to differentiate the source of specific sensor data. The increase in granularity, theoretically, allows for applications to be deployed at different levels in the naming hierarchy. [26] Suggests that current solutions in ICN naming are far from a leading naming solution. Although no work has proposed a universally adoptable standard, the fact that names are application dependent should simplify the global naming methodology adopted. The very nature of ICN, the liberty in naming and naming network opacity make it so that adopted conventions defined can be kept minimal. Furthermore, when contrasted with other fields in ICN research such as security or caching, naming is comparatively more advanced with matured works and surveys available in general ICN literature [29, 30]. The current proliferation of naming proposals for ICN content is associated with application specific optimizations and protocol dependent solutions, both of which are context specific. The heavily contextualized nature of solutions should limit the potential impact of individual solutions in standardization, with current works serving as reference rather than proto-standards. Lastly, In Chapter 5 we introduce the SEVeN architecture. As part of the proposed architecture, we introduce a prefix-naming solution that is based on application classification. This, as is discussed in the Chapter, was necessary to allow for more efficient Quality of Service management.

3.2.2 Routing and Forwarding

With an increase in promising applications or VANETs, the associated delivery mechanisms for each individual application have become a point of contention in VANETs. [52] present a classification for delivery solutions in VANET, classifying works based on replica allocation and routing model. In general terms, the approaches taken to content routing in ICN are classified as either Reactive, Proactive or Opportunistic. Within these approaches, many authors have proposed a variety of routing algorithms uniquely designed to tackle specific problems relating to VANETs. [53] discusses content discovery in MANETs. The work denotes the limitations in resource availability, service discovery and content provisioning. While this work is outside the scope of ICN, the discussion made has been used by many as a starting point for their work. In [54], the authors evaluate the efficiency of NDN in vehicular networks through a prototype test platform. The work evaluates two test applications, simulating requests for text and multimedia-based applications. Both simulations and real-world testing denoted the benefits of ICN, removing the isolation between different requests for the same content. The results obtained denotes the longevity a content object cached in the network can have, outliving its original producer. Similarly, [55] and [56] two nearly identical implementations of NDN for VANETs based on IEEE 802.11p, replacing the TCP/IP MAC layer managed by IEEE 1609.3 with NDN while retaining WSMP as its communication protocol. Content retrieval in the proposed solutions is divided into provider discovery and chunk retrieval. Performance evaluation in both works done using very simplistic Manhattan grids, but denote improvement of load distribution away from a single RSU to the vehicles in the network. The authors also discuss chunk-level content retrieval and its implementation in 802.11p. [57] introduced the first Interest flow control protocol for NDN based on an AIMD mechanism. The proposed algorithm utilizes a sliding-window model to regulate the Interest propagation at the receiver. Evaluation of the model assumed a high skewness in the frequency distribution and allowed the observation of NDN transport in caching. From this, the authors provide an optimal model validated by an analytical framework that ensures fairness. [50] introduces a routing solution named ICoC, explicitly designed to improve QoE of multimedia streaming services. It utilizes data mules and RSUs to improve content delivery. Data mules are classified as either partners or couriers depending on their location and lane information. Partner vehicles are defined as those within the same lane within a one-hop communication range, and couriers are those not in the same lane of the producer are labelled as couriers. The differentiation in classification is necessary given the potential changes in lane speed and are used for hop-selection during the routing process. The cooperative approach is designed to ensure that playback buffers at the client side are kept filled. [58] proposes a set of changes to the transport-level functions of ICN in vehicular networks including content segmentation and interest retransmission. Interest propagation and chunk sizes are modified based on network load, with positive results. [46] introduces a proactive routing protocol for ICN VANETs bloom filters. Entitled Hierarchical Bloom-Filter Routing (HBFR), the solution is designed to serve both shareable and non-shareable applications. The bloom-filter approach requires that urban maps be organized into geographical partitions and that vehicles be grouped into clusters. At the cost of additional overhead compared to the traditional NDN approach, HBFR decreases response time and traffic for unpopular content objects. This decrease is a consequence of the pseudo-geo-routing scheme that is implemented by the architecture. The application also requires that service providers be aware of whether the content is shareable or not, with no clear definition of how this is implemented being presented. Lastly, the estimation of content popularity is based on previous local statistics without information exchange, which in real-world scenarios are likely to reduce the performance of the approach. [59] introduces VENDNET, yet another variation of NDN designed for vehicular networks. The goal of this work is to implement communication schemes specific to the different scenarios in VANETs, namely V2V, V2I and V2R (Vehicle-to-RSU). The evaluation scenario considered is highly specific, and the system was designed to operate using LTE mobile networks. The network architecture is structured based on a static group of cells to which vehicles travelling along straight lines connect. The limitations of the simulated scenario significantly restrict the validity of the results obtained. [60] also studies interest propagation in ICN-VANETs based on NDN. The authors analyze the effect of different properties in VANET efficiency including vehicle speeds, hop distance to content and network sizes. This short work denotes the negative impacts of increased distance to content and the need for efficient interest propagation strategies. In [61], the authors introduce RUFS, a forwarding selection algorithm for ICN-VANETS. The approach is based on proactive satisfied interest statistics sharing, which is then used for interest forwarding selection. The solution limits interest flooding based on content satisfaction probability and distance to expected providers. [51] proposes a distributed mobility management scheme based on the content names and vehicle mobility information. The architecture introduced performs packet-aggregation based on mobility. Mobility information is also used as part of the routing algorithm, which restricts forwarding based on vehicle location. [62] introduces a geo-based forwarding strategy for NDN, modifying interest and data packets to include vehicle position information. Vehicles tasked with forwarding content explicitly compute forwarding lists. [63] introduces the LER geo-routing protocol for ICN-VANETs using social knowledge and last encounter information. This approach is designed to reduce content flooding and opportunistically route content. The proposed solution significantly reduces overheads from flooding and increases content delivery. Its greatest limitation is the need for location data to be shared. [64] presents another geo-routing protocol ICN-VANETs. This approach takes into consideration vehicle location and trajectory when determining its routing algorithm. [65] propose a lightweight interest-broadcasting protocol to mitigate broadcast storms. The proposed protocol considers the distance between a current forwarder and its neighboring vehicles to opportunistically limit Interest packet transmission redundancy. [66] presents LOCOS, a location-based content distribution protocol. Lastly, [67] introduces a geo-routing protocol for ICN-VANETS that modifies data names to bind the provider position to data object names. This work intersects routing and caching for applications that have explicit location dependency. The resulting adaptive content discovery mechanism efficiently identifies the best data sources based on the information propagated as part of the content request.

3.2.3 Caching

The intrinsic applicability of caching in vehicular networks offers opportunities in many avenues of research. This includes the evaluation of existing policies as well as the introduction of novel solutions targeted specifically towards the unique environment of VANETs. [68] reviews a series of classic cache solutions with applicability in ICN. The authors present and discuss multiple in-network caching solution, all of which are applicable within ICN VANETs. [69] evaluates cache invalidation schemes for temporal and location-dependent requests in mobile environments. While this work does not consider ICN as its scope, the scheme evaluated and focus on mobile make it a good reference point. Moreover, the consideration made on location dependency is extremely relevant for ICN, as many applications are likely to display this behaviour. The location models analyzed were characterized using cell-based location encoding and different level of validity information. The authors demonstrated the trade-off between performance and scalability for different degrees of precision of validity information attached to objects. Performance analysis of caching policies in NDN is presented in [70]. The authors introduce a caching algorithm based on node centrality. They also denote the relative efficiency of random replacement when compared to the pervasive caching approach adopted by NDN. Through both analytical and simulation-based evaluation, they demonstrate that both the proposed solution as well as LRU-based random replacement outperformed ubiquitous caching. While the proposed solution is not directly applicable in highly mobile VANETs, both the concept of centrality-based caching and discussion that denotes the issues with the "always cache" approach of NDN are relevant. [71] introduces the Cache Aware Target IdenTification (CATT) cache policy for ICN. In the proposed model, internal application and file-specific policies determine whether content can be cached. The theoretical benefit of CATT is that internal and external caching decisions are made at the time content is located. The content provider, based on the path towards the requester, determines the caching decision for the return path. While CATT is interesting, the ever-changing nature of vehicular networks significantly restricts the applicability of this approach, which requires relatively high knowledge of network structure to function properly. In [72], the authors introduce the ProbCache Policy for ICN. The policy takes into account the distance between the original content provider and its current position in the return path to determine the probability of it being selected for caching. The goal of this approach is to disseminate copies throughout the network while avoiding over-replicating copies. [73] performs a series of analysis of off-path content caching. The authors evaluate three hash-routing techniques, comparing their performance to the probabilistic cache approach. All three approaches outperform ProbCache and denote the benefits of increased dispersion of cached content throughout the network. Unfortunately, the simulation scenario observed is static and highly specific. Nevertheless, the work demonstrates the benefits of dispersion-based off-path caching in ICN. In [74], the authors criticize the performance of LRU-based cache replacement policy and introduce two alternate approaches. The proposed solutions are designed to optimize content replacement either only as part of the return path or on a network-wide scale. Interestingly, the authors consider a large selection of parameters to define their Mixed Integer Program (MIP) models. The two models presented are dependent on knowledge of content popularity, which significantly reduces the applicability of this approach in the real world. Nevertheless, the works presented serve as a point-of-reference for the potential cache-replacement algorithms may achieve. [75] evaluates the benefits of ICN in VANETs in terms of two routing characteristics: source selection and caching. From their analysis, they denote the benefits of pervasive caching and suggest shortest-path routing as the preferred solution for ICN-VANETs. Unfortunately, the authors perform their evaluation on a simple grid scenario and only evaluate the basic "always" cache policy against one where content can only be stored at either recipient or source. Nevertheless, they denote the benefits of caching in ICN. As part of its cooperative routing solution, the work in [50] introduces a collaborative caching mechanism. Partner vehicles that accept assistance requests download/cache contents from RSUs for themselves and their neighbour vehicles. This process is triggered by special Interest requests labelled for partner assistance. [62] introduces a small set of heuristics to reduce redundancy in cached copies in NDN architectures for VANETs. The proposed solution is based on defining a subset of vehicles to be cache insertion targets rather than having indiscriminate caching. In [59], the authors discuss Prefix-Tree (PT) sharing as a mechanism to estimate popularity based on prefix frequency, denoting its improvement over basic LRU-based replacement. However, the scenario selected is very limiting, which reduces the relevance of the performance analysis. Supplementary discussion on the topic of caching in ICN is presented in Chapter 6. Additionally, two novel solutions introduced in this work are described1 in Chapters 7 and 8.

3.3 Concluding Remarks

In this Chapter, we presented the current trends in research relating to ICN-VANETS, related work and the main contributions in the field. From the body of work summarized, it can be seen that there are still many open research challenges in ICN and its applicability in vehicular networks. Among the research opportunities listed, we pinpointed the lack of architectural components necessary to implement service exchange systems in VANETs. Similarly, we pinpointed limitations in currently existing caching solutions that can be applied to the vehicular environment. Many of the problems that still need to be tackled are derived from the unique challenges faced by vehicular networks, and, consequently, we understand that there is a need for additional dedicated research on the topic. From these conclusions, in the following Chapters, we proceed to introduce novel concepts that we believe mitigate some of the limitations discussed.

Chapter 4

Public Transportation in Vehicular Networks

The structure of this Chapter is the following. Section 4.1, contains brief introduction and contextualization of the theme of public transit in vehicular networks. Following, in Section 4.2 we discuss the general characteristics of public transit and describe the technology deployed to the location used for this case study. Then, in Section 4.3 we discuss the potential roles public transit systems can have in vehicular networks as well as challenges that need to be overcome. In Section 4.4 we supplement this discussion by analyzing a real-world trace of public transit mobility to the network coverage that can be achieved by PTSs. The Chapter ends with closing remarks, presented in Section 4.5 .

4.1 Introduction

The near pervasive presence of public transit systems, coupled with the very nature of the "public services" they provide, denote the critical stake that public mass transit has in the future vehicular networks. Public transit systems serve as the primary mode of transport for large portions of the populous and are for the most part are reliable and accessible. The public transit ecosystem, while being a fundamental component of the vehicular environment, however, does not its due research focus. With vehicular networks coming to fruition as part of Intelligent Transportation Systems (ITS), it is increasingly important to understand the potential roles of public transit. Public Transit Systems (PTS), despite their relevance, are infrequently discussed in the context of vehicular networks. Public transit is sidelined in the existing literature, only being considered in the contexts of data-muling and routing in proposed solutions [17]. This can be partially attributed to the limited knowledge on the topic of public transit. One of the reasons behind the lack of research on this topic is the challenge of obtaining and evaluating public transit data traces. Current analysis focuses on the mobility patterns of users rather than on the PTS itself. The limited roles associated with PTSs in the current state of the art are restrictive. We argue that a shift in vehicular network research is necessary to allow for public transit to become more central players in vehicular networks. From this context, this Chapter pursues the goal of increasing the understanding of public transportation, how it can be integrated into future solutions and what challenges must be tackled. We focus on correlating properties commonly observed in public transit systems with potential applications as well as challenges associated with the theorized applications.

Towards the goal of understanding the potential roles and challenges associated with Public Transportation System integration into Vehicular Networks and IoT, for the remainder of this Chapter, we study public transportation from the perspective of its users and service providers. As part of our evaluation, we analyze comprehensive data-sets provided by OC Transpo, the public transit authority of Ottawa, Canada.

4.2 Public Transportation Technology

Public Transportation Systems (PTS), or mass transit systems, are publicly available services used by billions worldwide as a means of locomotion. PTSs are provided either directly by local authorities or by private companies under concessions awarded by municipalities. The modes of public transit deployed to urban regions are rail and road, with bus service being the option that is more frequently available. Increases in the incorporation of integrated electronic fare-collection systems into public transit systems, alongside increases in wireless communication and location tracking integration, have significantly increased the research potential of mass transit and, consequently, its role in vehicular networks. Given increases in technology deployment to both private vehicles and PTS systems, it is natural to consider the potential roles that mass transit will have in future urban connectivity landscapes.

By integrating public transportation into their ITS architecture, cities can significantly reduce costs in infrastructure deployment with the added benefit of increased flexibility and upgradeability [76, 36, 35]. The two main benefits we see in PTS-ITS integration are *flexibility* and *cost reduction*. The control over PTS by public authorities means that deployed vehicles can operate as part of a trusted environment. The nature of public-service also allows PTSs to operate either as authenticators or mediators in applications that require trust management. Their presence as public service providers allows PTSs to act, for example, as access points, complementing road-side units (RSU) and reducing the need for fixed infrastructure deployment. PTS systems can be complementary to infrastructure deployment, covering areas that lack Base Station coverage, as well as assisting in strategic locations such as high traffic areas. The mobile nature of PTS systems also allows for strategic deployment and reassignment, with reduced re-deployment costs when compared with installation of additional infrastructure.

Bus-based public transportation systems are among the most frequent and widespread PTS systems available across the globe. The financial constraints associated with PTS system deployment tend to favor adoption of solutions with lower implementation and maintenance costs, preferring bus services. Buses do not require significant infrastructure to function, as they can share the road with cars. Restricting service deployment to a single platform reduces the variety of elements that require maintenance and improves maintenance efficiency. Being solely dependent on bus service is not detrimental to the quality of a city's public transit, as buses provide a flexible service. Buses are prevalent because fleet deployment is flexible both in terms of route assignment and fleet deployment. This flexibility allows for schedule changes to be made on-the-fly based on specific needs such as temporary events such as congestion, weather or social gatherings. Still, even if buses have a nearly ubiquitous presence, PTS are still restricted limited by the volume and capillarity of its fleet, routes and by the limitations of the technologies deployed.

Unlike private vehicles, which are not subject to extensive governmental oversight, public transit systems operate on budgets and are expected to avoid frivolous spending. Accountability and financial scrutiny significantly hinder public service operators potential to innovate by incorporating novel technologies into their systems. While installing a set of sensors in a vehicle might not be unreasonably expensive and provide a useful feature, deploying a specific technology to a fleet of thousands of buses can become prohibitively expensive. Hence, while academia and industry may introduce myriads of technologies, only those with financial incentives are likely to be adopted. Technology deployment in PTSs is informed by either the direct benefits brought by their inclusion or to remedy social outcry. One clear example where conflicting solutions have been proposed relates to inter-vehicle communication systems. The technologies most likely to be implemented for Vehicle-to-vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communication systems are *5G* or *DSRC* [37, 77]. Both approaches offer unique advantages and may have support from groups of manufacturers an governments. In practical terms, however, support for either technology will depend on the perceived utility and associated costs. The specific needs of each public transit operator will inform which technology is selected and it is possible for different regions to adhere to different standards.

Government oversight and financial restrictions, however, do not imply that mass transit systems are outdated and in dire need of improvements. While public services are frequently associated with poor quality, many cities are serviced by highly efficient providers that deploy modern technology to their fleets. One such example of a PTS operator with a modern fleet is that of the Ottawa region in Canada. This region is served by a single operator named OC Transpo, which employs buses as its primary mode of public transportation. All buses in the fleet are equipped with a modern set of hardware components. The equipment includes 4G LTE mobile network communication subsystems, onboard storage, GPS location tracking, fare tracking, passenger counting, digital sign panels, speaker notification systems as well as internal and external video surveillance cameras. All deployed subsystems are integrated into a single In-vehicle IP network, which produces, stores and forwards the information to a central Transit Operations Control Center. Buses have onboard storage with upwards of 10 Terabytes of data, 2 Terabytes of which is occupied by the video footage corresponding to the last 7 days, all of which can be accessed by the Control Center. In-vehicle sensor data and GPS information are pooled and forwarded to the control center at 30 second intervals. Buses consume on average $2GB$ of data every month for with a daily average of $5Mb$ of non-video data being generated. Collecting this data is relevant to the operator as it enables OC Transpo to preemptively remove buses from circulation for maintenance based on sensor information as well as reacting to issues observed while the buses are circulating. Hence, public transit already has the potential to provide many services solely with the technology that currently deployed. Following we present and discuss the roles that public transit can have as part of vehicular networks and intelligent transportation systems based on the introduction of new technology.

4.3 Roles and Applications

So far we have introduced and contextualized public transit bus systems. As different PTSs can facilitate distinct sets of services and individual applications can benefit specific subsets of features provided by PTS, it is important to understand the mass transit can have when integrated into vehicular networks. In this section, we present the potential roles and applications of public transport as part of vehicular networks.

Naturally, buses should take same roles as other members of VANETs, necessary for them to communicate and be part of network-wide services. In this context, public transit participation in vehicular networks does not predicate distinction between public service providers and private vehicles, though policies can be put in place in specific scenarios. Buses can also operate at a higher level of privilege than other vehicles, as they possess greater hardware capabilities and are public service providers. Following, we discuss a series of applications and how PTS can benefit from PTS VANET integration.

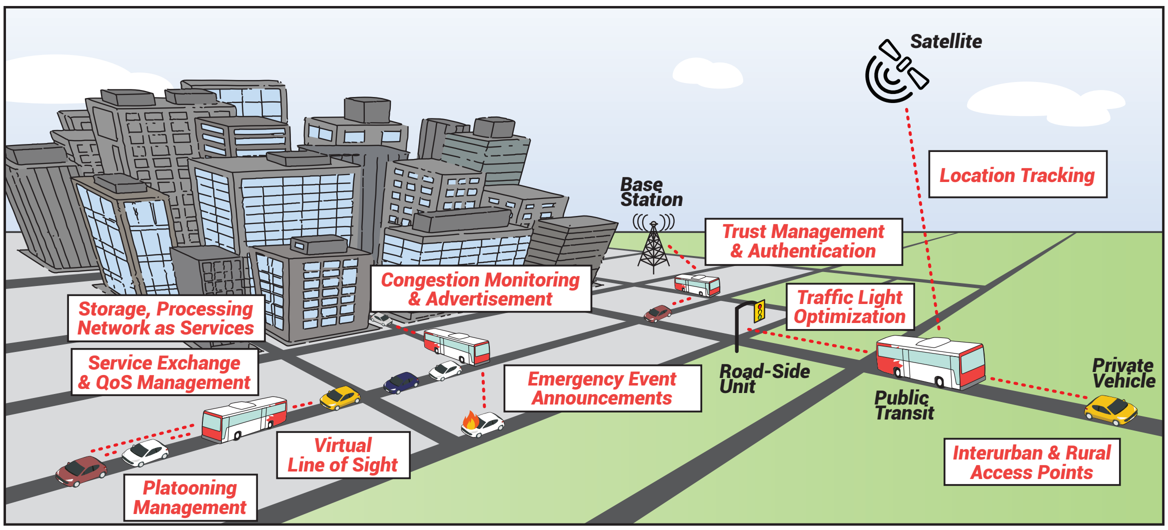

Figure 4.1: Mass Transit Integration into Vehicular Network Example Scenario

4.3.1 Service Exchange

The ever-increasing volume of sensors deployed to vehicles has significant potential to enable knowledge exchange services. Vehicular networks were initially proposed as a mechanism for emergency and safety warning dissemination. However, increased technology deployment allows for a vastly more extensive range of services to be introduced. The variety of contexts includes the proposal of applications outside those safety-related, and include traffic information dissemination and infotainment services [78]. Competition for network resources is expected to become an issue, a near ubiquitous challenge faced by mobile networks [79]. The increased competition will make network resource management a future necessity, as discussed in Chapter 2. Service and network access management remain open challenges that can be tackled by Public Transit. PTS nodes operating as official entities can, for example, be tasked with resource and service management. Ensuring fairness in access to network resources and services is fundamental both to manage competition and avoid predatory behavior but also as a mechanism to increase the trust level in the service exchange process. Increased service management capacity can stimulate participation in knowledge and service exchange schemes by curbing malicious and greedy behaviors. The participation of PTSs in an official management capacity would allow for increased coordination efficiency and subsequently improved service exchange quality. Should network resources become significantly scarce, access to network resources could be further managed by PTSs. PTSs can assist provide prioritized access to vehicles participating in services that benefit the community. The increased capillarity of public-based services provided by cooperation between PTS and infrastructure services can increase the efficiency of such management schemes. This increase can lead to the introduction of more complex and ambitious resource management solutions. Participation in applications relating to a public utility, altruistic knowledge dissemination or other services that benefit the network as a whole could warrant benefits for respective nodes. Sharma et al., for example, propose a Blockchain-based resource sharing model in which service requests operate hierarchically [80]. While their solution includes a management scheme operated solely by infrastructure systems, their approach can be expanded to also include PTs systems as part of the management model. This example speaks to the relevance and applicability of PTS integration in VANETs as service orchestrators. Other than managerial attributions, PTS vehicles can actively participate as regular members of vehicular networks, collecting and disseminating information or by providing services to private vehicles. Participating as a client in information dissemination services, naturally, is beneficial to public transit operators, as their vehicles can gain access to knowledge which can be used to improve service efficiency.

4.3.2 In-vehicle Passenger Services

Commuters can significantly benefit from PTS VANET services both in terms of network access as well as regarding specialized services. Given that commuters are passive actors in PTS transportation, significantly increases their probability of consuming content during their trip.

Passengers have potentially increased connectivity during their travel by accessing the Internet using services provided by public transit. This increase in connectivity is beneficial as it reduces contention for the already scarce mobile network resources while allowing for the PTS node to manage resource allocation themselves. While passengers can independently use mobile networks to access the Internet, PTS connectivity intermediation is useful as the proximity to the network access point allow for increased power control, potentially reducing energy consumption, essential for mobile devices. The connectivity capability of PTS vehicles is expectedly higher than that of regular users, enabling buses to achieve better quality and, in some cases, allow for connectivity in instances where individual users would not have any access to network resources. Underground systems and rural bus system are two examples where such events are likely to occur. In both cases, PTS can improve the quality of services delivered to users without the need for additional infrastructure.

Buses can also provide services other than connectivity. Public transit can relay mobility information obtained from the vehicular network to keep its users aware of potential delays, emergency related events, and other relevant information about their trip and the region including weather patterns and expected delays. Non-essential information sharing can also be enabled in PTS systems with digital media as well as advertising being made available for consumption. Increased service consumption can also drive monetization models, which can gate access to multimedia content or Internet access. Monetization can offset the costs of maintenance of technology or even provide additional funding for PTS managers. Advertisement content can also be made to take advantage of social characteristics of passengers. This type of dynamic advertising can be significantly more efficient than print media as content would not be physically attached to buses, enabling more directed targeting. Credentials associated with public transit could also be used as gating systems, enabling increased personalization of advertisement content. The degree to which PTS systems will be explored for intra-vehicle service provisioning will develop based on its revenue generation potential. In the case of multimedia content delivery, the storage potential of buses also allows for large volumes of content to be stored and made available to passengers without the need for connection to the Internet. Both Storage as a Service (SaaS) and local streaming services can be incorporated to provide entertainment systems for users. In-flight Entertainment System technology already exists in the aviation industry, and it is increasingly common for airlines to offer locally stored content to passengers accessible for consumption with their devices.

4.3.3 Delay Tolerant and Opportunistic Networking

The hardware capacity of buses can also be explored to enable implementation of specialized network services that require trust and dedicated storage. Buses can be tasked with collecting data from Delay-tolerant Network (DTN) devices. Predictable schedule and communication capacity are two driving factors in enabling PTS systems to serve data collection and distribution services. The long-term regularity of PTS route scheduling allows for planned collaborations with devices that lack an Internet connection. Large volumes of data can be collected en-route and delivered as buses return to their terminals. Mesh network sinks and other distributed networks can take advantage of this predictability to cyclically forward data to PTS nodes. The fact that public transit systems are accountable further assures potential partners that offloaded content will be delivered. It also allows them to deploy additional physical hardware buses for data collection purposes. Suburban and rural area networks such as crop, cattle, and wildlife monitoring systems can benefit from access to PTS provided connectivity. Alternatively, these clients can gain intermittent access to the Internet via intermediated infrastructure access relay. PTS systems can operate as opportunistic network access points in suburban and rural areas covered by transit authorities. DTN has seen significant research in the context of MANETs and VANETs with works focusing on platooning, geo-routing and sensor data collection. While the related literature contains research specific to data-collection and transmission using PTSs, existing solutions can be expanded to be integrated into general VANET systems rather than being deployed as one-off services.

4.3.4 Authentication

Generally considered an untrusted environment, vehicular networks are susceptible to a wide variety of attacks, ranging from impersonation and repudiation to congestion and denial of service. Implementing trust is fundamental not only as a general communication requirement but also to enable many of the applications that have been envisioned for VANETs. The nature of mass transit allows it to operate as trusted and authentication environments, which can significantly assist in securing and validating vehicular communication. Buses can operate as edge nodes, forwarding requests to infrastructure tasked with authentication, extending the range of trusted environments. In scenarios where infrastructure-based authentication is required, PTS nodes can operate as edge devices, intermediating connection between individual vehicles and the infrastructure services. Alternatively, PTS systems can themselves serve as authenticators. In the case of independent identity validation, PTS nodes can, without the need for infrastructure, operate cell-based identity management schemes. By requiring proof-of-identity such as License plates or other registration information, PTS nodes mitigate threats that arise from vehicles falsifying their identity and can improve privacy. PTS systems can also assist with privacy by allowing nodes to anonymize their identities. By maintaining association tables, anonymous aliases lists can be kept for each user, allowing neighboring vehicles to communicate without the need for unnecessary information disclosure. As discussed previously, buses are already capable of exchanging information in real time with PTS control centers and already collect passenger data during boarding. Lastly, buses can also operate as monitoring and enforcement agents for infrastructure services, monitoring and collecting data. Potential applications include identity validation and impersonation detection via sensor or other side-channel analysis. As an example, buses could use camera sensors or directional antennas to detect malicious vehicles attempting to impersonate other vehicles.

4.3.5 Challenges

Implementing additional services into already existing Public Transportation Systems is not without its issues. Successful deployment of services and integration of PTSs into vehicular networks requires that the many hurdles be overcome. There are many fundamental challenges associated with the implementation of PTS-based services. Given that VANETs are still in theoretical stages of research, proposed solutions need to be flexible to adapt to an evolving playing field faced with ever-changing technology and increasing contention for network resources and services. Maintaining flexibility is further complicated by the bureaucratic nature of public service. Naturally, individual PTS applications will have specific requirements depending on the application-space in which they participate. Nevertheless, there are general hardware and software requirements that will need to be faced.

From a technical standpoint, the requirements for successful PTS VANET integration can be divided into prerequisites shared with other vehicles belonging to VANETs and unique requirements necessary for the implementation of specific roles attributed to PTSs. If currently deployed hardware serves as a point of reference, we can safely assume that technology deployment is only a limitation if there is no political interest. Given efficiency or financial incentives for their inclusion, it is not unthinkable to predict that buses will meet all the requirements to provide the services listed previously. Nevertheless, given the need for accountability, public transit should not be expected to be early adopter of any given technology. This way, unless governmental mandates are enforced, such as the case for the DSRC standard [81], public transit systems are not likely to be early technology adopters. The need for interoperability further denotes the challenges associated with technology adoption. Given the potential for public transit to operate as a mediator between private vehicles and the infrastructure, interoperability becomes paramount in public transit vehicular network integration attempts. Ultimately, public transit is likely to adopt whatever solution is necessary to enable compatibility with the remainder of the vehicular network, meaning that PTS systems will likely select the de-facto standards of their region. As far as dedicated equipment requirements, to minimize expenditure, we predict that additional technology deployed will only be introduced to provide specialized services.

The most impacting non-technical issues are associated with policy making. Private interests, lobbying, and budget limitations can introduce significant delays or restrict technology deployment for public service providers. Even in the face of significant benefits, initial deployment costs can lead government entities to reject projects designed to introduce new technologies. This way, the degree of independence of public service operators and their ability to justify allocating resources is a determining factor in the viability of application deployment. Reduced hardware requirements and standardization of components should lessen the need for additional spending and, consequently, expedite service deployment processes. From this, we conclude that is fundamental to understand the potential benefits of PTS, including its current context and potential, to justify the integration of PTS systems. Following, we perform a study using real-world data to denote the potential for bus-based network connectivity services that are aimed to encourage additional research on the topic.

4.4 Connectivity Study

Mass transit systems, unlike traditional infrastructure-based network service providers, offer networked services using mobile interfaces. This way, the quality and usefulness of network services provided by a PTS are heavily dependent on service availability. As PTS-based network services are to be provided by mobile devices, the expected network coverage and service quality fluctuate as the vehicles travel. Hence, PTS-based network services should not be expected that have a comparable quality of service to those provided by infrastructure. Nevertheless, it is essential to understand the capacity PTS has of providing networking services to potential clients.

For this evaluation, we explore a comprehensive dataset, made available by OC Transpo, which contains one full week of GPS tracking for its entire fleet of over 900 vehicles. We evaluate bus mobility to derive the potential network coverage provided by Ottawa's bus system. Each record in the dataset contains the identifier, time-stamp and last known location for a bus. With this dataset, we aim to estimate how much of the road networks would have service coverage provided only by the bus fleet.

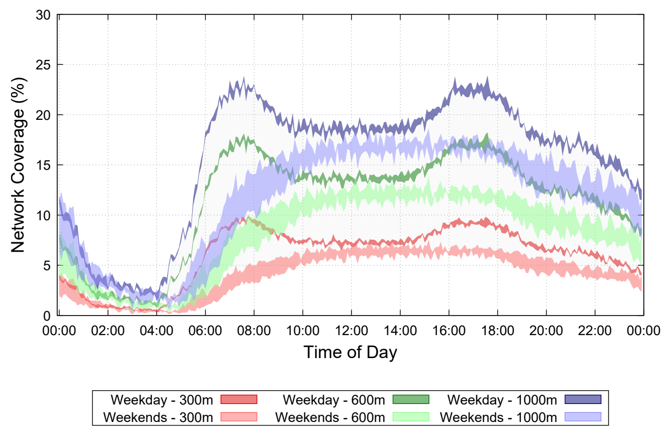

(a) Network Coverage over Time for Otawa's Central Area

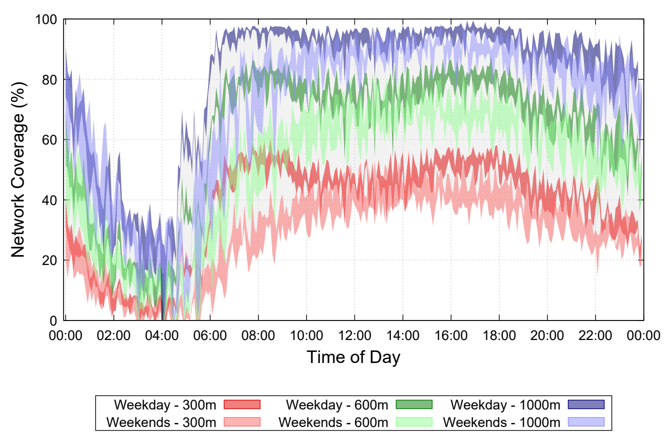

(b) Network Coverage over Time for Otawa's Central Area

Figure 4.2: Average Coverage Distribution Over Time

The approach taken is based on measurement of the road coverage by buses over time. We divide the dataset into daily $24h$ cycles further divided into time slots with 1-minute duration each. For each time-slot, we compute the last known position for each bus actively running a bus route based on their path and their last known location. Buses that have not reported any updates in the last 5 minutes are considered to have completed their routes and are not included in the analysis. From the list of GPS locations, we build a representation of road coverage based on radial coverage without restricting connectivity ranges based on potential interference sources such as buildings. The resulting list of maps each contains a series of R radius discs each representing the area covered by a bus actively operating a route overlapped over the physical representation of the road network. The resulting maps are used to evaluate the portions of road surface that has network coverage. For each time-slot, we calculate the resulting network coverage. To better understand the effects of asymmetries in bus routes along the region observed we analyze both the estimated coverage for the entire region as well as an analysis of the only the roads contained within the city's central area. As we have access to multiple days worth of tracking data, we present our results as ranges defined between the minimum and maximum observed values for each scenario configuration analyzed. Additionally, considering previous observations on the clear distinctions in mobility patterns between weekdays and weekends, we have also included separate measurements for days belonging to these two groups. Lastly, to account for different potential technology deployments, we present separate measurements for three different coverage ranges: 300, 600 and 1000 meters. These values specify a range approximating the maximum distance attainable for IEEE 802.11p devices at default transmission power, estimated 5G connectivity range based on preliminary testing and estimated LTE-V2X range [82, 38, 41]. The resulting network coverage curves are presented in Figure 4.2.

Analyzing the results for the entire urban region, which are presented in Figure 4.2a, we see significant potential for bus-based service provisioning. The Ottawa urban area is vast, encompassing a total $5,7156\ Km^2$. Up to 23% of the region's road network is covered during business days and 18% coverage is achieved during weekends. For the more restrictive connectivity configuration, where the connectivity range is 300 meters, coverage during the business week is limited to 10%. From the results, it is clear that weekend coverage measurements are frequently lower than their weekday counterparts, due mainly to the reduced number of buses being in activity during similar times. This reduction is expected and is a consequence of the difference in the number of buses actively operating routes. The difference in coverage between weekday and weekend statistics is most significant during the early morning rush hour, which takes place between $6AM$ and $8AM$. This difference is another consequence of changes in scheduling, with a temporary increase in bus deployment during weekday rush-hours to compensate for the increase in demand, behavior which is not observed during weekends. Moreover, the coverage disparity is not linearly proportional to the difference in the number of buses actively operating routes, as during business-days there is increased coverage overlap, especially in the central core of Ottawa's downtown. Nevertheless, during extended business hours ($6AM$ to $8PM$) network coverage is reasonable even in worst-case scenarios. The fluctuation between the minimum and maximum coverages does not correlate linearly to the difference in ranges and is a result of coverage overlap. Buses are not distributed evenly along the region with routes converging to the city's central area, which leads to a significant overlap in this region. The increases in network coverage observed during rush-hours, another common characteristic of public transit is also impacted by network overlap. We observed an increase in the number of buses active between 100% and 200%. However, the increase in network availability was disproportionately small, with relative increases between 25% and 35%. This increase is due to how buses are distributed along space, with significant coverage overlap in the central region. This area has increased population and business density, and it is also where vehicle traffic is most dense. Consequently, it is relevant to evaluate this area on its own. Statistics specific to the central region only evaluation In Figure 4.2b. As can be observed, network coverage is significantly higher in Ottawa's center area. While coverage is expected to be higher than the region average, during business hours, network coverage is never lower than 40% with nearly complete network coverage being achieved during all business hours for the 1000 meter scenario. These results are promising and denote the significant potential that buses have as service providers and participants in vehicular networks.

4.5 Concluding Remarks

In this Chapter, we performed an in-depth evaluation and discussion on the characteristics of public transit and its potential role as part of vehicular networks. We contextualized public transit as a pivotal yet neglected topic and presented a series of contexts in which their participation can yield network-wide positive impact. Then, we theorized a series of applications that bus services integrated into and roles it can take as part of in such groups. Ranging from service provisioning to information exchange and authentication, we presented the benefits of including buses into vehicular networks. We also summarized the biggest challenges that need to be overcome in order to allow for PTS integration. We presented issues specific to each of the roles previously discussed and how they relate to current research and technology deployment. To evaluate potential service provisioning we also performed trace-based analysis, which denoted significant potential with over 20% city-wide road network coverage being achieved during a large portion of the work week and full coverage of central regions during most business hours. The discussion and the results presented in this chapter supplement the discussion presented in Chapters 6 and 3. Together, the content presented so far serves as a sound technical basis and are a reference point for those interested in ICN-VANETs. In the following Chapters we build on this knowledge to introduce novel architectures and protocols.

Chapter 5

SEVeN: A Novel Service-Based Architecture for Information-Centric Vehicular Network

This Chapter is structured as follows. Section 5.1 provides the context that motivates the introdution of SEVeN, the novel architecture presented in this Chapter. In Section 5.2, we present the relevant characteristics of Information-centric networking considered in our work. Open topics in ICN and related literature of ICN-VANETS are presented in Section 5.3. Following, we present the SEVeN architecture in Section 5.4. Section 5.5 presents a simulation-based performance evaluation of the proposed model. Lastly, we present closing remarks for this chapter in Section 5.6.

5.1 Introduction

The advent of near-ubiquitous connectivity is observable and poses significant challenges in areas such as Vehicular Networks. Vehicular Ad-hoc Networks (VANETs) were originally purpose-built systems designated to enable safety related systems such as emergency break notification. Nowadays, however, services such as video streaming represent a significant portion of network traffic and is projected to account for 82% of traffic in 2021. Moreover, Mobile to Mobile (M2M) traffic is predicted to grow more than sevenfold in the same period [14]. To cope with increased content demand and variety, VANETs design needs to be revisited. Moreover, unlike issues such as energy limitation, processing, and storage capabilities, connectivity remains a noteworthy limitation.

The challenges imposed by mobility, which further hinders communication, needs to be addressed and incorporated into vehicle-specific solutions.

The backbone of the Internet was designed based on the host-to-host communication (IP) paradigm. However, content distribution and caching have become significant features with communication moving towards content delivery [4]. Clients are interested in the content itself with source location not being relevant. An Information-Centric solution, in theory, can solve previous issues while enabling efficient management of the application ecosystem [83]. In ICN frameworks, the communication paradigm centers around content. Amongst the proposed architectures, Named-Data Networking has attracted significant research attention [26]. In NDN, *Interest* messages containing the name of the desired object are used for information retrieval, and the content names are used as part of the routing process. Rather than forwarding requests towards specific hosts, request lookup occurs on a per-node basis and compare content object names to local Content Stores. The premise of ICN is that by decoupling content provisioning from physical content location, the communication process becomes more flexible.

Vehicular Ad Hoc Networks were originally conceived to promote the safety of life [84]. Dedicated Short Range Communication (DSRC) systems were designed with specifically to enable safety related functionalities [81]. However, technology has evolved significantly since DSRC's initial inception in 1999, and research focus has broadened the potential application horizon for vehicular technology in terms both of device deployment as well as application variety. Sensors deployed to vehicles are a potential source of knowledge and have been extensively evaluated by existing literature. Ranging from image to speed and spatial awareness sensors, vehicles are increasingly capable of producing knowledge that can be shared amongst vehicles. These sensors are useful not only to their local vehicles but also to other members of a VANET. Information sharing also allows for improvement outside the safety-related scope, enabling local traffic knowledge sharing, as well as infotainment related activities to be carried out. Considering the increased application-space of VANET and expansions in infotainment service availability, VANETS are faced with increasing demand for services.

One of the main consequences of service availability and variety is the need for improved access management and content exchange schemes. It is natural for VANET services to prioritize applications relating to its original purpose as both need to ensure that safety requirements as well as adhering to legislation. However, this does not mean that service management should be made to exclude applications outside safety and transportation scopes. In fact, it is fundamental for additional services to be incorporated, albeit with restrictions. The use of network services in the context of multimedia and entertainment allow for greater adoption of VANET specific communication systems and for decreased dependency on infrastructure-based mobile networks, which are costly and already face significant demand. Moreover, as VANETs move away from safety-exclusive contexts to more general-purpose applications, service access management become a fundamental issue and one that has currently does not have a clear solution.

The increasing service availability and variety and limitations of physical addressing motivate the drive for more efficient communication. Given the limited network resources and the inability of IP-based solutions in VANET environments, we consider service-management as implemented by ICN networks as a viable alternative. Towards this goal, this Chapter presents Service-Exchange Vehicular Network - *(SEVeN)*, our novel service-oriented architecture for information-centric vehicular ad-hoc networks (ICN-VANETs). SEVeN implements a service architecture that extends and enhances features provided by NDN while enabling application co-existence and service exchange. By taking advantage of service classification and class-based prefixing, the SEVeN architecture allows class-based prioritization, ensuring that fundamental services such as those transit-related be fulfilled without application-specific prejudice to lower priority services such as infotainment applications. Based on SEVeN, we proposed the *Load-Based-Deprioritization (LBD)* content prioritization policy. Simulation results denoted the need for efficient service access management schemes as well as the increased delivery via use of the LBD deprioritization mechanism proposed for SEVeN, which increased both total delivery rate as well as prioritized category delivery rates with minimal changes to latency.

5.2 ICN Background

In this section, we describe the network access model considered as the basis for the work presented in this Chapter, the policies and other elements that are relevant to the architecture we propose. It has been shown that named-data content delivery can cope with a high level of mobility, maintaining high request reply rate unlike host-to-host networks, which cannot maintain comparable service quality [12]. The work in this Chapter furthers this research effort by proposing a system architecture based on information-centric networking. VANET applications have a wide application range with equally ample requirements. As discussed previously, the increasing volume of services motivates research in information-centric networking. More specifically, the limitations of IP in VANET environments, coupled with the special requirements of vehicular networks motivate research on service-oriented solutions.

Amongst the existing proposals, the most adopted ICN solution is the Content-Centric Network approach, which has been expanded and renamed as the Named-Data Network (NDN) model [9]. While alternative approaches to ICN such as DONA, PURSUIT and 4WARD exist, we consider the NDN architecture to provide the most flexible foundation for the implementation of ICN [23]. NDN is simple and provides a general framework, serving as basis for several works. This approach approach moves the universal component of the TCP/IP network stack into named data object chunks, as per the ICN paradigm. NDN is consumer-driven and functions based on interest propagation. Communication is local at each step so all sender-receiver pairs are individually involved in managing their network load. Communication is structured around two basic message types, Interest and Data packets. Interest Packets are used to broadcast requests, disseminated via flooding. Incoming interests are checked against the Content Store (CS) for the longest-matching prefix and returned if a match is found. Otherwise, the request is compared to the Pending Interest Table (PIT) and either added to a content interest list or discarded had it been previously included in said list. By having pending requests being stored in the PIT, broadcast storms are partially avoided, as subsequent requests for the same object are not rebroadcasted within the requests lifetime. This reduces the need for cycle and knot detection in routing protocols [85]. Interest requests are forwarded and incoming responses to forwarded requests are stored in the Forwarding Information Base (FIB), with outgoing interfaces connecting requesters to forwarders being registered for future reference. NDN interest packets also include nonces to prevent duplicate interests from different paths to be rebroadcast in a loop. Further improvements to content discovery have been published, ranging from reactive, to proactive and opportunistic coordination schemes [33].

Content naming policies are another fundamental feature of Information-Centric networks. In NDN, consistent global names are achieved with the use of a prefix-based naming. Names are hierarchically divided into components with arbitrary octet sizes. NDN does not specify any naming-conventions, with the responsibility over naming structure is determined by applications. Notably, it has been shown that there is a tight relationship between ICN namespaces and application contexts and that therefore prefix naming structures only need to follow a minimal set of conventions [86]. Hence, by employing a policy that structures information based on application context we can reduce prefix sizes and improve network efficiency lookup times. Prefix-based naming also allows for service grouping and service substitution. Services from alternate providers can be substituted or derived from more general categories via extrapolation. Prefix generalization can also potentially reduce network load by allowing for requests to be fulfilled by providers closer to the requesting node.

Strategies for on-path caching based on information prioritization, as well as security and name-based cryptography, need also be considered in the definition of namespaces. While existing works in VANETs for these research topics are limited, there are solutions for content caching available in ICN literature [73]. Similar considerations apply in regards to service discovery, discussed in depth in related literature [53, 46, 87]. General-purpose ICN solutions are divided into reactive, proactive and opportunistic. Reactive solutions are dependent on the propagation of interest requests, being flooding-based. Controlled flooding can potentially reduce the overhead caused by content request broadcasts and is achieved via distance, neighborhood and geo-aware solutions. The application of a specific reactive routing is dependent on the specific properties of applications. Traffic-related information, for example, is area specific and road-segment information is likely available at the location for which a request is made. Solutions, however, are dependent on application-specific properties and no single approach is universally capable of delivering improvements. Proactive discovery and routing allow for vehicles to advertise services and content. Notably, while proactive advertising necessarily increases overhead, reductions in the cost of service discovery can merit preemptive broadcasting. The granularity of the items divulged plays a significant role in determining the overhead and utility of preemptive sharing. Consequently, the viability of proactive advertising is dependent on the duration of connections and on the overhead result of broadcasts. Opportunistic solutions offer an alternative that mitigate the flooding overhead by implementing passive discovery mechanisms relying on agent, social and intelligence-assisted mechanisms at the cost of increased latency. Based on background characterization presented this section, we proceed to the description of the SEVeN system architecture, described following.

5.3 Related Work

Information-centric networking is at the forefront of innovation in their respective research fields. Therefore, while a topic of significant research interest, the volume, and scope of related literature is limited. Most works focus on system-level solutions for traffic information dissemination or on defining conceptual abstractions to classify ICN VANETs with contributions ranging from content classification schemes to network concepts [33]. Liu et al. classify works in ICN into two sequential stages, those focusing on theoretical issues and those focusing on application-specific solutions respectively. The first stage saw the proposal of many approaches to ICN including NDN, as discussed in the previous section. The second stage saw the introduction of many works, some of which we discuss following in this section.

Dedicated Research focus has been devoted to the evaluation of previously proposed system-level solutions and its components. The viability of named-data communication in VANETs is demonstrated via a simplified computer simulation and physical network tests in [48]. The authors propose a Name Data Network VANET model with multiple network interfaces, and a content sharing model based on different service classes. The evaluation process, while simple, classified content into one of two types of requests: Info-Traffic and Road-Photo, with packets emulating localized chunks of traffic information and geolocated traffic photos. Existing caching solutions are evaluated in the context of VANETs and ICN-VANETs [74].

Lee et al. [88] present a general discussion on open research topics in VANETS as well as challenges in vehicular Cloud Networking. The authors also propose several definitions on network structure and communication services. A very interesting point mentioned is the evaluation of incentives from a service perspective, where they mention the need for rewarding service provisioning. This notion is fundamental in systems with limited resources and a multitude of information sources. Recently, Chen et al. discussed service-oriented networks in terms of IoT, which further denotes the relevance of the approach [89]. However, to the best of our knowledge, no works have tackled service-oriented communication in the context of ICN-VANETs. We tackle this research shortfall with the introduction of SEVeN.

Most efforts, however, focus on context specific solutions [33]. Multimedia streaming is studied by Quan et al. [50]. Considering the limitations in networking resources, the authors propose a cooperative caching model to improve multimedia stream delivery for highways. The proposed solution takes into consideration vehicle arrival rate and relative linear vehicle motion and assumes small relative speed between vehicles. Similarly, several works discuss content caching strategies for location and application specific contexts [63, 46]. Only recently was content prioritization introduced as an alternative to access management [90]. The authors propose a back-off based approach to interest and data forwarding based on content popularity perception. While the work has limitations, including the need for global knowledge of packet popularity and limited capacity to cope with high network load it useful in demonstrating the need for efficient service management schemes.

Additional research efforts have been made in VANET Cloud platforms. Solution range from parking lot cluster-based cloud service management to interest-based service clouds [91, 92, 93, 94]. Two common themes discussed by works on this topic are the lack of coordination between service providers and limited network resources. The increased interest Information of Things (IoT) services and subsequent knowledge production is the basis for a data mining cloud system proposed in [91]. The authors argue that the lack of global standardization in device integration and communication architecture require further efforts to allow the implementation of coordinated vehicular data clouds. Reliability in service clouds is studied by [92]. The work proposes a controller-based access protocol to maximize network and average per-client access rates. A contextual interaction model that enables users to exchange content without resorting to the standard Client-Server model by using generic platforms is proposed in [93]. The work in question considers network function visualization and software-defined networking as fundamental components to achieve its goal. The authors adopt a REST API as the basis for their solution, which they argue allows for the maintenance of the connection between clients and servers regardless of interaction failures. The proposed approach, however, is limited to session caching and retransmission as means of maintaining connections. A Cloud platform system based on an extended RESTful system is then proposed by the same team in [94]. The authors focus on semantic-based services as a basic element for cloud formation as they discuss the benefits that semantically rich applications attain from higher-level views of their topologies.

Evaluation of the existing body of work in ICN VANET makes clear that the research area is still in early stages and that significant effort is required in several topics. Two of the main issues with existing solutions is the low level of abstraction of proposed works and the lack of validation of proposed solution. Current works in ICN VANETs limit their studies to either overall evaluations of the feasibility of ICN VANETs or contextually specific solutions that target a specific application. As denoted by Liu et. al [33], current trends of ICN suggest that most solutions focus on application specific scenarios that promote self-improvement of architectures. Moreover, considering ICN-related research is at a relatively early stage and that most works target very specific contexts, there is a natural interest in converging the different applications into a single connection model. Given the current research context and predicted application ecosystem to be implemented into VANETS motivate the proposal of a service-based architecture for ICN VANETs and definition of a system-level solution that allows for an integration of different services.

5.4 SEVeN: A Novel ICN Architecture for VANETs

In this Section, we present and discuss the system architecture of SEVeN. Our solution takes into consideration the explosive growth that Vehicular Networks have experienced in recent years and the increasing trend of connectivity. Trends such as the Internet of Things, Smart Cities and their effect in the implementation of Intelligent Transportation Systems (ITS) illustrate the belief that service provisioning and management needs to be a core feature of vehicular networks. In both instances, proliferation and variety of accessible services motivate the introduction of coordination systems such as the IFTTT [1] web-based smart-things automation service. Notably, however, the high level of mobility, limited coordination and existence of real-time applications motivate the introduction of a solution tailored for vehicular networks. Towards implementing a system that enables effective VANET communication, we introduce the SEVeN architecture, centered around service and knowledge exchange based on the NDN general model. We consider that by enabling managed content exchange that considers individual vehicle preference and content classification that significant improvement in content delivery can be achieved. The predicted benefits over physical addressing include improved content access, routing efficiency and caching, as described in Chapter 1.

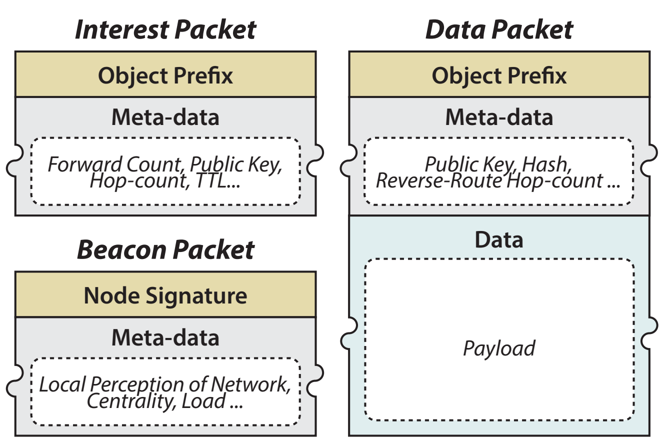

Figure 5.1: SEVeN Packet Types and their contents

[1] www.ifttt.com

SEVeN significantly expands on the core concept of NDN and is tailored specifically for VANETs, focusing its design around the peculiarities of the field. Given its VANET-oriented nature, SEVeN takes advantage of features present in the vehicular environment. Given regulatory mandate in the US by the FCC, we opted to the vehicle-specific IEEE 802.11p/1609 protocol family for vehicular environments, generally known as WAVE (Wireless access in Vehicular Environments) As per WAVE standards, we include beacon-messages into our communication scheme, expanding NDN's message incorporating beacons as a third message type. Beacon exchange in WAVE was designed as part of traffic safety application operations and is implemented via regular advertisement frames, mandatory for all vehicles in the network. Notably, the utility of having vehicles share additional information during beaconing regarding their perception of network state as well as other potentially relevant information is promising. As an example, beacon-assisted GPS coordinate exchange can be used to accurately track neighbor location, allowing for increased routing efficiency [95]. This way, we envision beacon messages being used for multiple purposes, expanding on their initial utility and allowing for increased network knowledge and, subsequently, efficiency [95, 96]. The *heartbeat* beacons are exchanged in a specialized Control Channel (CCH) while other packet types are broadcast in the four existing Service Channels (SCH). Significant research has been devoted to understanding beacon-casting and to improve beacon package reception. It has been concluded that with adaptive transmission and frame-fitting, that it is possible to efficiently manage beaconing and associated beacon-dependent applications [97, 98, 99]. Consequently, we are confident in the inclusion of beacons as an explicit packet within the SEVeN framework. The three basic packet types considered for SEVeN are visualized in Figure 5.1. Both Interest and Data packets, are centered around content exchange and have prefix-based data object name as their fundamental fields. All packet types include additional meta-data headers, which vary depending on the packet type in question. Regardless of type, packets are predicted to include additional information such as timers as well as elements specified by NDN. Despite an increase in the utility of beacon messages, beacons remain separate from other message types in SEVeN. Supporting evaluation and discussion is presented in Section 5.5.3.

Additional meta-data can be incorporated into the naming scheme using a modified naming format such as the one proposed by [50]. However, request prefixes should be minimized to reduce header and lookup time. This way, additional meta-information used to describe service requirement as well as additional filters and other meta-data objects separated from prefixes interest messages whenever possible. Applications that include a high level of granularity such as location specific information, should adhere to either grid-based location systems such as the MGRS (Military Grid Reference System) or a road segment name structure depending on their context as these models reduce interest-flooding by structuring location specific requests and therefore enabling directed interest forwarding [33]. Privacy can also benefit from the inclusion of meta-data in content requests. The incorporation of an asymmetric key algorithm such as the Public Key Infrastructure (PKI) system used by web services is a simple mechanism for user authentication. In this case, requester and content provider public keys can be included as meta-data along content request and reply packets to ensure authentication. Similarly, cache access and copy restrictions, authentication, insertion recommendation, service exchange, as well as other policies can also take advantage of meta-data.

5.4.1 Communication Scheme

The communication model proposed for SEVeN follows a requester-oriented interest propagation approach. In NDN, applications are responsible for naming their data and forwarding requests to the network given they adhere to the basic naming policy scheme. It is the responsibility of applications to specify their naming schemes to minimize interest prefixes. Notably, while content objects should be globally understood, meaning is only relevant at the service level and naming can be compacted to further reduce prefix sizes. Regardless of local understanding of the contents of a content object, all nodes should be capable of understanding any request and, moreover, its general purpose. In SEVeN, we propose an expansion of the concept of “global naming” to one where nodes can derive properties of a content object from its name. We consider it fundamental for vehicles to be able of perceive messages as useful, high-priority or even private to allow for increased awareness. Our proposal is based on the specific requirements of VANETs and its mobility related features and is implemented as part of our prefix model.

VANET network communication patterns are categorized based on the properties of the data requests. We consider the most commonly used categories, namely, *safety services*, *transit information* and *infotainment* in our work. The different content categories display varying characteristics, including popularity, network relevance, priority, Quality of Experience (QoE) and bandwidth requirements. The properties that define objects define patterns that we incorporate into the naming scheme of our solution. We consider prefixing-based classification, as exemplified in Figure 5.2. The prefix model proposed is designed to enable service management. Similar to most prefix-based systems it is

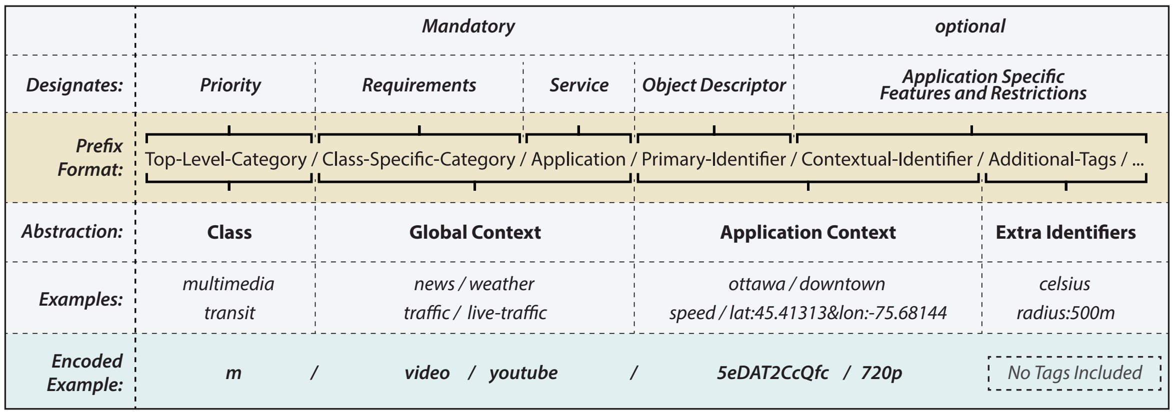

Figure 5.2: SEVeN Prefix-based Naming Structure with Examples

boundary-based, with '/' serving as the component separator. Unlike the more generic structure defined in NDN, the prefix system of SEVeN proposes a explicit hierarchical subdivision of elements centered around the root element. The root element of the prefix in SEVeN is labeled as the *Top Level Category* and represents one of the three main categories and to which the request belongs. The content types, which we refer to henceforth as a *content class*, have associated properties including *class priority* and purpose. Designing class description into the prefix model allows for content prioritization schemes to be implemented efficiently while making request purpose explicit and unambiguous. Consequently, global prioritization is applied to the content classes to avoid prejudice to specific services.

Applications, which we describe in terms of the services they offer, are grouped into *Class Categories*. Both safety and transit application classes are designed to augment the capacity of fundamental services and therefore have more explicit application spaces. General-purpose applications, in contrast, are grouped into the infotainment class. Notably, class-specific is mandated, as to enable categorization regardless of application requirements. Class categorization is fundamental if one wishes to enable fairness within non-fundamental applications. By specifying application category as the second prefix component, policies that affect specific services within the infotainment scope can be made to enable QoS within a class scope. Individual categories represent a certain type of service, generally associated with the descriptive properties of the service described previously. While Quality of Service is applied at the top level category level to ensure that service priority is maintained, QoE requirements denoted by service categories are perceived based on class categories. This way, different applications from a same category, such as video services, are grouped into the same service category and are not segregated in terms of prioritization. *Application* level prefixes specify individual services, either relating to specific service provider types (e.g.: Youtube) or service categories (e.g.: Weather). This distinction is made based on the perception that sensor-based applications are not grouped based on the service provider name but rather on the features made available by them. The inversion between organization name and service category made for SEVeN in contrast to that of NDN is also based on the properties of VANETs. VANETs were originally purpose built and segregate applications based on their purpose, hence, making class specific categorization more relevant to routing than application names. The increased number of prefix fragments are not considered to increase lookup times as both top level categories as well as class-specific categories are globally defined. By employing binary encoding and prioritizing smaller prefixes for popular application spaces, total prefix sizes can be significantly reduced, with Service-related sub-prefixes being responsible for the greatest portion of unique identifiers. Moreover, the prefix structure can increase network efficiency in terms of policy execution, as only the initial prefix fragments may need to be interpreted during policy evaluation for interest requests. Lastly, it should be noted that similarly to NDN, named requests are self-certifiable via encryption and producer signatures. This way, system-level prefix obfuscation is fundamental, as managing access control in an ad-hoc system would introduce significant overhead with limited efficacy. Moreover, given the potential inclusion of privacy-related meta-data to requests, obfuscation brings little benefit. We note that anonymity and trust management, are two topics that remain open issues in ICN and for which we do not make additional considerations.

Services fulfill user demands for specific resources or communication needs, either physical (*e.g.: storage*) or logical (*e.g.:traffic data*). Consider the example relating to traffic information dissemination presented in Figure 5.3. Individually, any given vehicle will have partial knowledge of traffic conditions, still, via content exchange, all members of the VANET can become aware of full route transit information. In scenarios such as this, vehicles will take both roles simultaneously, sharing their traffic knowledge while requesting information from surrounding vehicles. Individually considered, information shared by vehicles is partial and at times conflicting, additional content identifiers and meta-information can be used to limit inconsistencies. In the example given, when a vehicle requests traffic information, they include region specifying filters as part of the

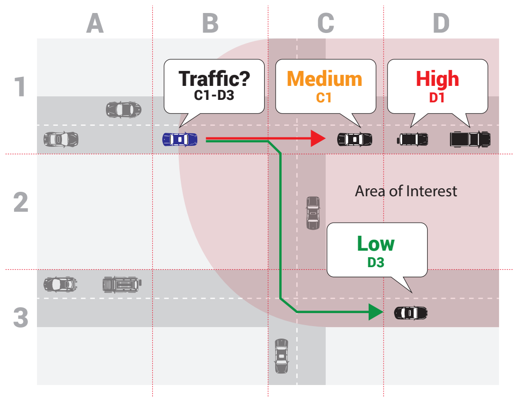

(a) Example of Information Area of Interest

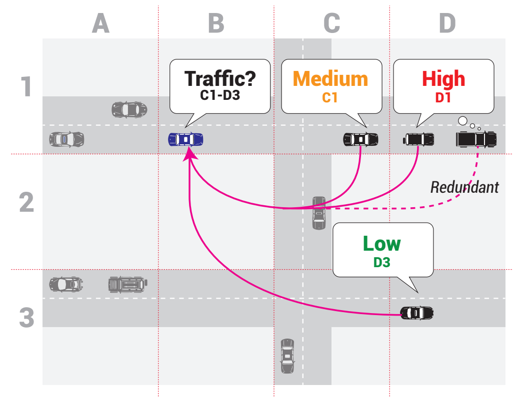

(b) Example of Information Redundancy

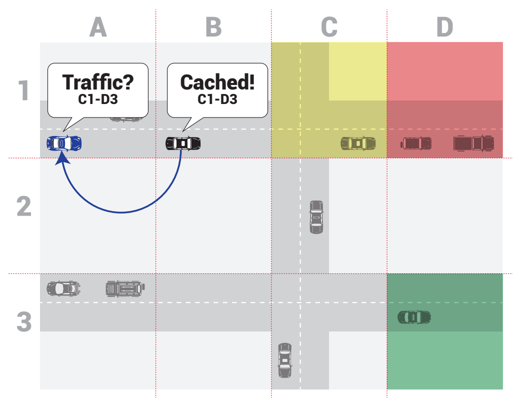

(c) Example of Information Caching

Figure 5.3: Application Example: Transit Traffic Status Requests

requests *Contextual Identifier*. By including the contextual identifier, prefix filtering leads to a reduced number of responses relating specifically to the requested area of interest. In fact, by employing network coding and response suppression mechanisms, overhead can be further reduced. We do not specify the properties of the area of interest for an item as it is dependent on the application context, similarly to other properties such as cache validity. Similarly, the location-specific property of traffic information can be taken advantage with the use of directed routing. Notably, as observed in the above example, information structure, cache policy, and other requirements are application specific. The varying application requirements make maintaining content within the content store a complex task. To cope with requirement variability we propose the *Service Sublayer*, which includes specialized components responsible for service management, discussed in

Section 5.4.2. In Synthesis, content names are structured around prefix-based service classification. Content classes have general characteristics associated with them, such as priority, popularity, average packet sizes. These properties are considered in the definition of the routing policies and are managed by the service management component of SEVeN.

All transmissions in SEVeN are broadcasts and Regarding content discovery and interface management, due to high mobility, preemptive content source discovery is not an attractive approach. Similarly to ICN, we forgo this approach in favor of requester-based reactive content discovery, augmented by opportunistic and social schemes. At this stage we only consider the inclusion of a single broadcast wireless interface in the form of a 802.11p interface, the FIB has limited applicability. Consequently, we do not make special considerations to the implementation and maintenance of the FIB. However, given foreseeable multi-homing result of inclusion of infrastructure base mobile communication technologies such as 5G, we do not explicitly discard the FIB, as it can be especially useful for applications that may require offloading. Any remaining characteristic relevant to the the communication scheme and that has not been discussed is inherited from proposed RFC drafts from the Information-Centric Networking Research Group (ICNRG) [2]. Next, we introduce the *Service Sublayer*, its subcomponents, each tasked with specific roles and related policies.

5.4.2 Service Sublayer

The choice to move towards an information-centric approach was based on the potential benefits of named-data in service exchange. By using location-independent data, SEVeN enables a migration to global service-based network structure with increased support for decentralized management. Furthermore, named-data allows for in-network caching, replication as well as content prioritization and content exchange as services. The system architecture of SEVeN is displayed in Figure 5.4. As can be observed from the diagram, our solution is designed to be integrated directly atop the IEEE WAVE Architecture, with transversal integration between the data and management planes as a mean of enabling inclusive security and management features as part of service provisioning mechanism. Specifically, SEVeN considers the basic components of NDN and includes a *Service Sublayer* as part of its variant thin-waist model, designed to provide service management and provisioning functionalities. The Service sublayer is designed to enable service management and information sharing schemes. Its primary roles are maintain-

[2] Drafts available at: https://datatracker.ietf.org/group/icnrg/

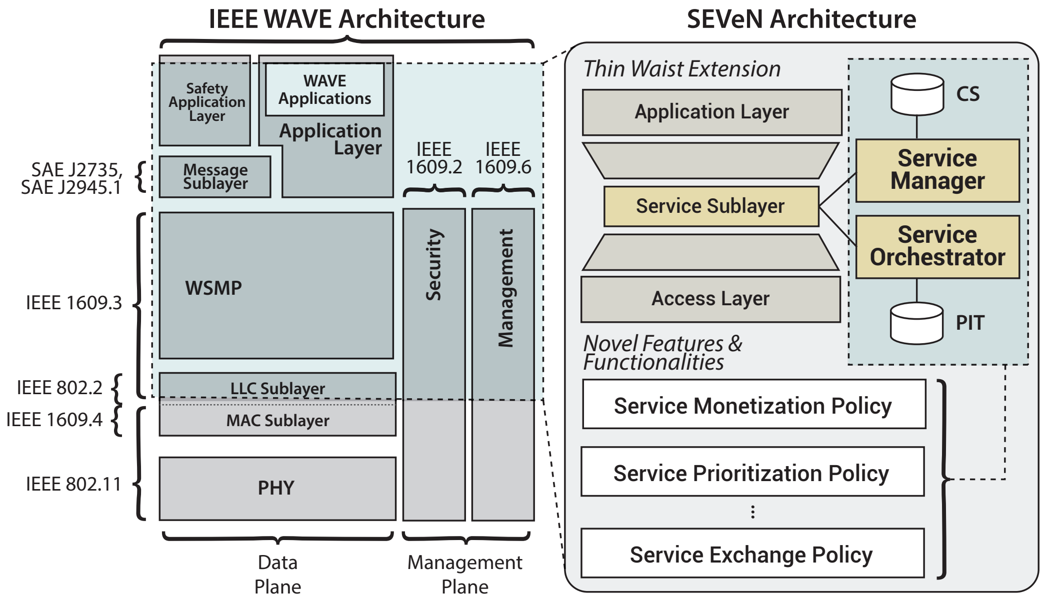

Figure 5.4: SEVeN System Model with Thin-waist Service Sublayer and Integration with IEEE WAVE

ing content caches, managing, and aggregate application data, broker service exchange with remote nodes whenever required and orchestrate coordinated network access. The functionalities proposed for the service sublayer are performed by two component, the *Service Manager (SM)* and the *Service Orchestrator*. The *Service Manager* is responsible for supervising service providers, collecting data from content producers and maintaining the content store. The *Service Orchestrator* is responsible managing network access and request prioritization. All nodes in an SEVeN network can behave as client and service providers, brokering and allowing access to services, which are provided by the service manager. The possibility of increased integration, service exchange and coordination between nodes is one of the many benefits of ICN over IP. Intermediate caching, for example, can be greatly expanded with nodes offering access to content after serving as intermediate nodes in previous communication cycles. Thus, by introducing the service sublayer, SEVeN leverages the potential in VANETs and enables different applications and other VANET use-cases to co-exist in a functional way. Additionally, the Service Sublayer participates in content forwarding processes and is responsible for intermediate caching. This way, adhering to NDN conventions, the service sublayer maintains its basic data structures (CS, PIT, FIB), which are incorporated into the service sublayer.

Note that due to its limited use in current VANET context the FIB is omitted from the diagram but remains a component in the SEVeN architecture. It should be noted these features are managed by either the Service Manager and Service Orchestrator.

5.4.2.1 Service Management

The primary roles of the Service Manager are obtaining and providing content and consequently maintaining the content cache, including on-route intermediate caching services To perform these roles, the Service Manager interacts with local service providers and the content orchestrator, meaning that it participates in service access mediation. Service Providers include sensors, databases and other sources of useful information. Local service providers inform their service manager and iteratively forward any knowledge generated in the form of s named-data objects. Similarly, local clients forward their content requests to the service manager, who evaluates interest against the content store. Notably, for services where content is not cacheable, the Service Managers evaluates the interest against the list of local service providers. Potential clients include applications, remote nodes or embedded sensors.

As named-data allows for increased cache management, the Service Manager is constantly active. Service provisioning to content from remote sources is coordinated with the Service Orchestrator. Nevertheless, the Service Manager is active in all service requests. Additional requirements can arise from certain content categories, potentially increasing service manager load. The Service Manager, while focused on maintaining the content store can manage individual application requirements to further filter and maintain its cache. Requests for content such as Traffic information have cache validity dependent not only by temporal freshness but also by other factors such as the presence of traffic accidents. In the example relating to traffic information, a network-wide events traffic incident reports would invalidate objects otherwise considered content. In the case of location dependent content, upon being transmitted away from its region of interest, factors such as temporal and location freshness will affect its cacheability. It is the responsibility of applications to define specific properties that affect cache management by the Service Manager. The reliability of the traffic data, for example, might be lowered to the point of invalidity depending how its cache policy is specified by its traffic management service. This example denotes the importance of cache control, and further denote the benefits of Information-Centric approaches, which can be enhanced by utilizing contextual information to improve the quality of network knowledge. Naturally, is unfeasible for a single cache management scheme to efficiently manage a diverse content cache. Towards enabling efficient caching, SEVeN considers cache control as a meta-service provided by the Service Sublayer. Basic caching is implemented by the Service Manager and additional cache policies are performed based on explicit specifications relating to the local validity, explicit lifetime and local interest of nodes as defined by applications [88]. The description of such properties is shared alongside the content itself as part of the meta-data header component of Data Packets. Trust management is not explicitly listed as a service manager as content is self-certifiable. Hence, applications should consider the assumption that they operate in an environment without trust and, therefore, are responsible for trust management.

Content request frequency is another feature that affects the efficiency of the Service Manager. As content delivery rate and time are strongly affected by the average distance between request client and content source [57], network efficiency is improved by popularity based content caching. This is one of the motivations behind the inclusion of *Beacon Messages* as part of the SEVeN architecture. While no central entity is present, nodes can take advantage of beacons to exchange information regarding request frequency and type. By allowing vehicles to exchange basic access pattern descriptions, cache coordination can be produced, even if in a limited manner as coordination is ultimately dependent on sojourn times. Access frequency exchange allows for coordination and knowledge dependent solutions, such as the priority models, to become viable.

Similarly, as the cache should operate at line-rate, distributed cache coordination is limited [29]. This way, the standard operation for SEVeN nodes is that of full independence with passive implicit coordination being achieved using beacon messages. Ultimately, this means that the caching policy has limited knowledge of network state and that cache policies need be kept simple. The service manager is restricted by these limitations when maintaining the cache. However, these restrictions are not a significant hindrance to the effectiveness of the service manager. This is because it has been shown that simple cache policies produce satisfactory performance results [70]. Related works have demonstrated that random replacement algorithms attain performance similar to that of *Least Recently Used (LRU)*. While specific policies have been designed to take advantage of spatio-temporal properties, popularity or other data, for general purposes, LRU is a fully functional and simple solution. This way, we consider LRU the default cache cache eviction policy, with all incoming requests being eligible for cache. Should a content class or object not be eligible for cache insertion, that requirement is made explicit via the inclusion of a meta-data field for non-cacheability.

Lastly, the Service Manager is responsible for Local service provider coordination. While many services are dependent on network communication, access to storage, computing as well as other resources available are managed by the Service Manager. Hence, the Service Manager functions similarly to a resource scheduler, with control over varied service portfolios available locally and as part of content exchange mechanisms coordinated by the Service Orchestrated. Both modules are to operate in constant communication, coordinating access to content locally or remotely otherwise.

5.4.2.2 Service Orchestration

While the Service Manager functions as a librarian and operates similarly to a run-time infrastructure, the Service Orchestrator takes on the role of a negotiator. The reasoning behind the existence of the Service Orchestrator is the need for resource management, a result of limited connectivity. The basic features performed by the Service Orchestrator are service access management and brokerage. In face of limitations in networking capacity and overwhelming opportunity for knowledge sharing, we consider the implementation of service exchange systems a fundamental component of a service-oriented architecture. Hence, we consider a service brokerage a practical solution for scarce resource management. Certain services have universal appeal and usefulness while other have a limited interest within the network. Given varying interest, a natural question for service-based systems is whether unpopular or high bandwidth consuming services should be deprioritized or not. By implementing a brokerage-based service-exchange model, popular services are intrinsically prioritized in the network. Increase request frequency allows for greater probabilities of content availability at intermediate nodes. However, in face of high competition for network access, intrinsic service popularity is not enough to ensure expedient content delivery. To tailor content delivery to user interest, we consider the brokerage system of SEVeN to consider knowledge as a commodity. High-interest categories are not only intrinsically more relevant but they are preferable whenever resource competition is a necessity. This way, information sharing is itself can be interpreted as a service. Naturally, user individual interests can be fulfilled with node service orchestrators explicitly arranging with their peers for content to be provided. While we do not delve into the intricacies of brokerage policies in this work, we discuss the basic properties expected of individual policies. As discussed previously, services have associated measurable requirements, which can be explored in defining of pricing policies. By pricing the associated *cost* of accessing content available locally, nodes implement an information economy. One node, for example, can price its participation as an intermediate node during a series of communication cycles in order to gain access to low-priority content. The greatest challenges in solving this issue are the definition of pricing schemes and maintenance of distributed currency systems. For more information on these topics, one can refer to works on incentive-based distributed brokering services. For the sake of completion, we limit the considerations in service brokerage to the definition of service priority, described in the following Section.

The Service Orchestrator is also tasked with communicating with the entities responsible for network interface management. This entails the possibility of multi-homed devices to be controlled by a network management controller. While such task is outside the scope of SEVeN, the inclusion of a separate component specifically designated to integrate interfaces is relevant for many reasons. Technologies such as 5G can be relied upon for backup connectivity to the Internet should content not be available within the network or enable connectivity with legacy IP systems. As mentioned previously, we consider connectivity with infrastructure services a last-resort measure due to financial associated with mobile network communication. Moreover, data flows can be easily multiplexed without application awareness, and as DSRC predicts communication system interconnectivity, this consideration is not infeasible [33]. In fact, the fragmentation in VANET deployment, result of varied specifications being suggested by different research groups suggests that future vehicles will be multi-homed, participating in multiple networks simultaneously. However, as VANETS are still in a relative infancy stage, this discussion is not expanded.

Lastly, the Service Orchestrator is responsible for classification-based content prioritization. By exploiting data prefixes, the network can choose to prioritize relevant information. This task is fundamental considering the need to ensure QoS for priority top-level classes of content. While fundamental services operate at priority level and have dedicated channel allocation, non-emergency services need to share the communication environment. Access contention is a common pitfall of wireless communication systems, with unbounded access and exponential back-off policies significantly degradation system performance. While access control and interference avoidance are features ubiquitous to wireless systems, service classification allows for significantly improved access control. Hence, the Service Orchestrator inherits responsibility for NDN's PIT. By managing spectrum access to services, service-orchestrator can improve service quality to high-priority services and, whenever possible, satisfy the requirements of low-priority services. As per DSRC specifications, a multitude of public safety and transit management features are predicted to be incorporated. Considering the predicted increase in service variety with the incorporation of infotainment services, access orchestration is fundamental.

5.4.3 Service Prioritization

Service prioritization is based on the concept of “Knowledge as a Commodity”. As discussed previously, content popularity affects service usage, but in face of high levels of competition for network access, we consider it fundamental that service prioritization schemes be implemented. One mechanism that has been proposed for prioritization in the literature tackled this problem from a popularity perspective [90]. In the related work, the authors propose a popularity based prioritization scheme that employs a timer-based rebroadcast mechanism to delay low-popularity interest and data broadcasts. Their approach is limited by the need for global content popularity to be known by users. Moreover, the defer Window mechanism proposed has limited efficiency when network load is high, as will be discussed in the following section.

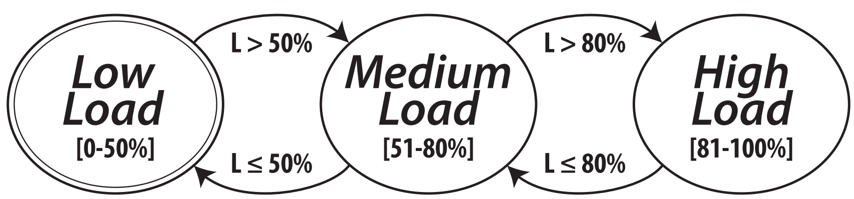

Figure 5.5: LBD System States and Transition Conditions

The implementation of an efficient is dependent on system capacity to maintain prioritization in high-load scenarios. To let high-priority packets be prioritized over low-priority ones, define a load-based state-machine presented in Figure 5.5. Our solution, dubbed Load-Based-Deprioritization (LBD), considers the definition of three explicit network states representing perceived *Low*, *Medium* and *High* network loads. Based on the state perceived by each node, they deprioritize packets depending on the top-level category of incoming packets (Transit, Network, Infotainment). Packet deprioritization is performed on top-level categories to avoid service-specific denial of service and comply with the proposed SEVeN model. Nodes do not have any a priori knowledge of network load, and load estimation is computed locally at each vehicle. Load measurement is performed with packet tracking that measures network use in bytes over a period of time. The relative occupation of the network is perceived based on the moving average obtained from past samples. The window size is estimated based on the exponentially weighted running average from the past λ samples, with each sample representing the bandwidth usage for sequential 1 second intervals. Network load values are obtained from the 1609.4 MAC layer based on the observed network occupation from both local node and remote transmissions relative to the transmission capacity of the network interface given in percentages. Occupancy statistics are sent to the service sublayer and are stored in the moving average sliding window. At every interval, the average load value is computed as $A_t = \alpha \cdot o_t + (1-\alpha) \cdot A_{t-1}$, where A_t is the average occupancy at instant t, o_t is the observed occupancy for instant t and α is the weight factor, set to $\alpha = 0.25 \cdot \lambda$. Based on the perceived network occupation, the service orchestrator transitions between states. Note that transition between *low* and *high* load states necessarily require for nodes to operate in a *medium* load state for at least one window slot interval. This decision was made as given the distributed nature of network estimation and exponential weights of the sliding window. The network policy adopts measures based on its current state. For Low-Load states, the LBD policy does not filter data or interest packets for any policy, forwarding packets without prejudice. For Medium-load states, LBD randomly discards $\mathcal{P}_{\%}$ of low-priority (Infotainment) interest packets. For High-load states, LBD randomly discards $\mathcal{P}_{\%}$ of medium-priority (Network) interest packets and all low-priority (Infotainment) packets.

5.5 Experimental Performance Evaluation of SEVeN

The focus of this Chapter is the definition of the SEVeN architecture, designed to enable service-oriented ICN-VANETS. Hence, to denote the benefits of SEVeN, in this section, we perform a set of simulation evaluations. Based on SEVeN's system architecture, we focus on evaluating the performance of its mechanisms, such as its prioritization policy in contextualized scenarios that denote the issues pinpointed by the approach. To evaluate the performance, extensive simulation was carried out using the OMNeT++ Discrete Event Simulator, Vehicle in Network Simulation (Veins) framework and Simulation of Urban MObility (SUMO) road traffic simulator setup [100, 101, 102]. The following subsections describe the properties of the simulation setup in terms of how traffic and content models were structured to enable the required evaluation.

5.5.1 Simulation Setup

The reference simulation scenario considered consists of the Ottawa's Sandy-Hill and Downtown Area. The simulated scenario occupies an area of approximately $2.5km^2$ and approximate linear road surface length of $29km$ distributed between *Residential*, *Secondary* and *Tertiary* road segment types. This scenario was selected given its grid-like structure and due to local knowledge of the region, fundamental in validating simulated vehicle behavior to ensure high fidelity level.

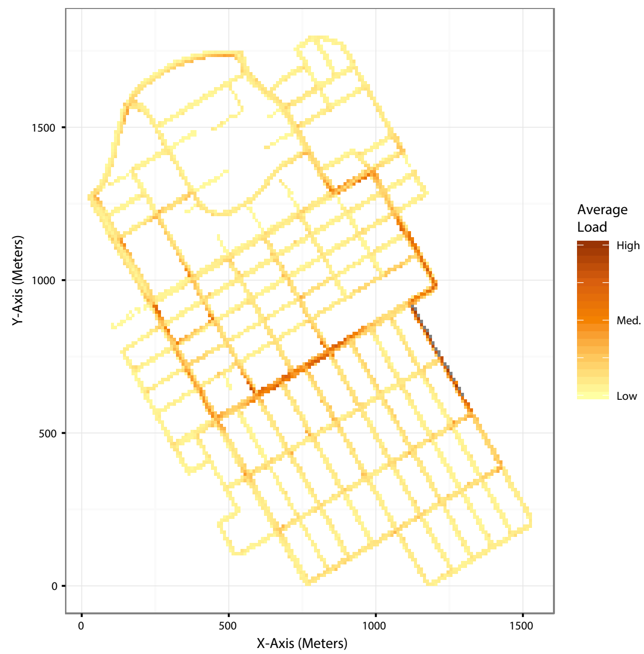

Figure 5.6: Heat map of Downtown Ottawa for High Density Traffic Scenario Configuration used for Traffic Flow Validation

Traffic conditions were determined based on road-type and real-world statistics and traffic flow parameters were obtained from real world areas. Traffic data with 24-hour period hourly averages from Edmonton and Calgary were used as basis to determine average vehicle flow rates [103, 104]. This way, despite traffic flow information not being available for the simulation area selected, traffic flow information from similar locations was used as basis. Average downtown road segment flow ranged between $3000 - 50000$ vehicles and hour-segment traffic flows ranging between $2.0\% - 7.2\%$ of total flow in the referenced material. From these parameters, we opted for vehicle density configurations ranging between $50 - 450$ vehicles. The density range selected represents traffic flows ranging from very low to high. The simulated scenarios are equivalent to vehicle densities between $2 - 20$ vehicles/km. Simulation results are presented for network scenarios for $\{50, 100, 150, 300, 450\}$ vehicles. Traffic densities and distribution were fixated for each density, validated according to predicted movement patterns observed in day-to-day observations of the region. The traffic densities are used as basic seeds for simulations with fixed test periods of time of 600 seconds (10 minutes) preceded by 100 seconds vehicle instantiation and startup times. The initial setup time was included to allow for vehicle instantiation and dispersion. Simulation density was kept content during simulation and movement patterns were verified using heat maps, as visible in Figure 5.6. A minimum 30 simulation repetitions were run for of each simulation configuration variations. The results displayed denote a 95% confidence interval.

The proposed SEVeN architecture and the NDN model it is based in were implemented into the tool-set according to the specifications listed in Sections 5.2 and 5.4. SEVeN is deployed as *clean-slate* solution atop the 802.11p framework available from the tool-set using Wave Short-Message Protocol (WSMP). The implementation of additional components required for the operation of our tests was done based on the specifications listed in Section 5.4.1. The ample nature of the SEVeN and its relationship with WAVE and NDN make so that certain details of the implementation have been omitted as they adhere to the technical standards and architectures upon which SEVeN is designed. In order to evaluate the effects of competition for network access, we generate artificial content objects to represent different object types similarly the approach in Grassi et al. [48]. For the content frequency evaluation, we consider two categories based on the global content classes: *traffic information* and *multimedia content*. For the prioritization evaluation, we consider the three basic content classes. Individual requests for contents follow the proposed prefix scheme but with only the Primary object identifiers being differentiated for each content object. Associated with each SEVeN vehicle we attached an emulated client that produced content requests. The requests were generated with a normal distribution based frequency and were associated with content classes based on frequency distribution described in the following sections.

5.5.2 Access Contention Evaluation

In this subsection, we evaluate the effects of competition for network resources by services with different bandwidth footprints within an unbounded scenario. The purpose of this evaluation is to further illustrate the limitations in network resources and need for network prioritization. Traffic information objects were configured to be $20kB$ in size while Multimedia objects were set to $5MB$ each. Network traffic content requests were implemented so that vehicles always forgo local knowledge in favor of outsourcing information regarding road conditions to their peers. Multimedia objects distributed evenly throughout the network, with individual caches being populated by a global library containing one thousand (1000) separate individual multimedia objects. The number of content holders for each multimedia object is approximately 10% at the start of every simulation. Consequently, content object availability maintains a fixed proportion with decreased distance between vehicles and potential content providers. Each client included a request timer and cyclically requested content to its SEVeN service manager. The category interest request was selected using a weighted random distribution based on the frequency distribution of the content categories. The proportion between the two request categories was varied from 20:1 to 1:4 traffic requests for each multimedia request. Maximum message payload was set to 1460 kilobytes.

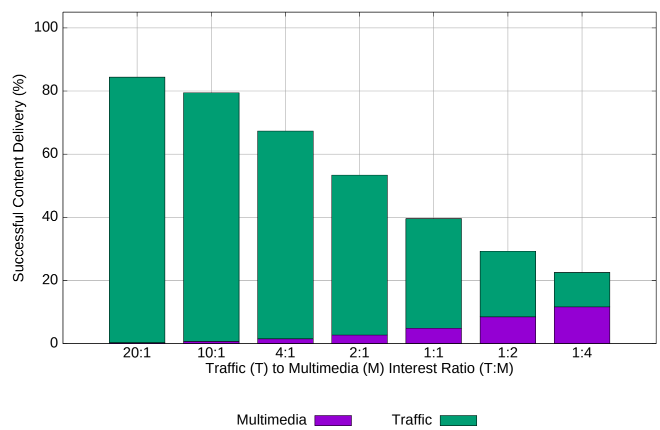

Figure 5.7: Class specific Transmission Success Rate for 300 Vehicle Traffic Density Scenario

Communication was measured in terms of successful packet delivery rate for both content classes. Summarized results for a sample scenario is displayed in Figure 5.7. The x-axis displays results for the different ratio between traffic and multimedia requests. For each ratio value, the y-axis displays the average percentage of successfully completed requests for each content category, with the aggregate from both categories varying between 0 and 100%. As can be observed in the graph, increases in multimedia request frequency

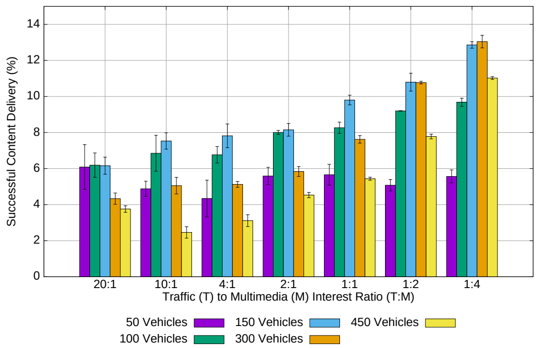

(a) Delivery Rate of Multimedia Objects

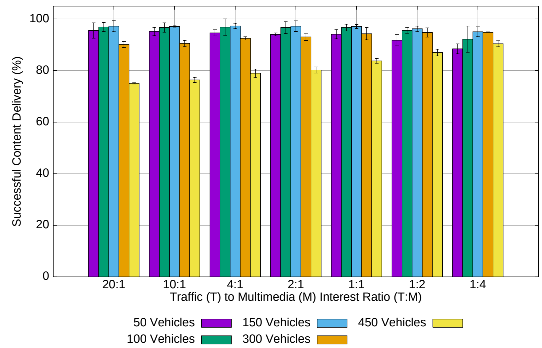

(b) Delivery Rate of Text-like Traffic Objects

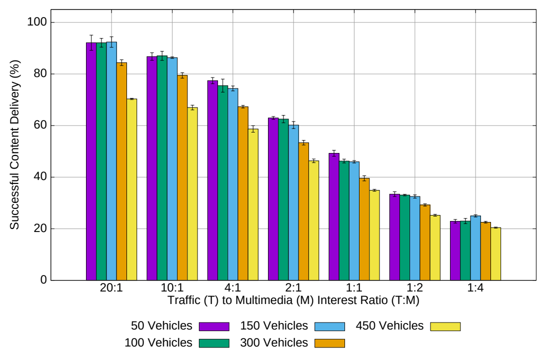

(c) Aggregated Results from Multimedia and Traffic Requests

Figure 5.8: Content Delivery Rates for Varied Traffic Densities for two different content types

reduce the total success rate of the system. This is a consequence of greater network occupation and subsequently increased packet collision. As larger packet requests become more frequent, overall delivery is reduced. Notably, the statistical prevalence of successful traffic transfers is replaced by failures in multimedia transfer. This trend is furthered detailed by evaluation of the delivery rates, as presented in Figure 5.8, which displays total delivery rates for both content categories for each scenarios. Notably, the frequency of successful multimedia transmissions remains disproportionately small. The increases in delivery are result of increases in the number of vehicles in the network. These results, as displayed in Figure 5.8a, are a consequence of increased service provider availability and average distance reduction in content provisioning. Evaluation of the effect of vehicle density in multimedia communication leads to interesting observations. While the results for low-density vehicle networks remain similar, increases in vehicle density lead to increased multimedia content delivery. This is a consequence of the content distribution model defined for the experiment, designed to increase item replication proportionate to increases in network density. Moreover, the results for the densest network topology ($450 vehicles$) consistently displayed lower delivery rates than sparser configurations. We attribute this result to a saturation in network occupancy and increased content provider availability which, together, further increase competition for network resources, leading to increased packet collision rates and subsequent decreased delivery rates. From the results, it is clear that unrestricted request forwarding is not an adequate network resource management policy as it leads to resource exhaustion.

Interestingly, as observable in Figure 5.8b, traffic requests are mostly unaffected by increases in multimedia results with traffic delivery rates for the category remaining stable, if smaller in absolute number. In terms of the effect of network density, interestingly, simulation results denote a degradation of up to 20% in the delivery rate of traffic data objects as a function of increases in vehicle density, observable for the 20:1 total success rate, as displayed in Figure 5.8c. This is likely the result of natural increase in network load result higher vehicle density. Figure 5.8c further denotes a reduction of up to 5× in total delivery rate with network density having decreased effect in the results as the proportion of vehicles increased. Considering negative effects of vehicle density and multimedia request frequency along with the notion that vehicle density is estimated to be increased up to 50% before total road traffic collapse, significant degradation is expected for extremely high-density scenarios. As network conditions have already shown significant deterioration in the 450 vehicles scenario, we consider the experiments set presented a good example and answer to the question of whether there is real need for the implementation of network access management policies and request prioritization schemes. Given the clear need for service access contention, we proceed to the evaluation of the service prioritization policies including SEVeN in as proposed in Section 5.4.3.

5.5.3 Beacon Contention Evaluation

Additional tests were performed to investigate the effects of beaconing in the communication process should beacons take place in the SCH. In this experiment, vehicle density was set to 300 vehicles and beacon frequency was evaluate for frequencies between $0.1s$ and $1.0s$, with beacons being sent in SCHs as well as the CCH. Simulation parameters were configured as they were in the experiments presented in Section 5.5.4 and results are presented in Figure 5.9. Similarly to the previous experiment, increases in contention for

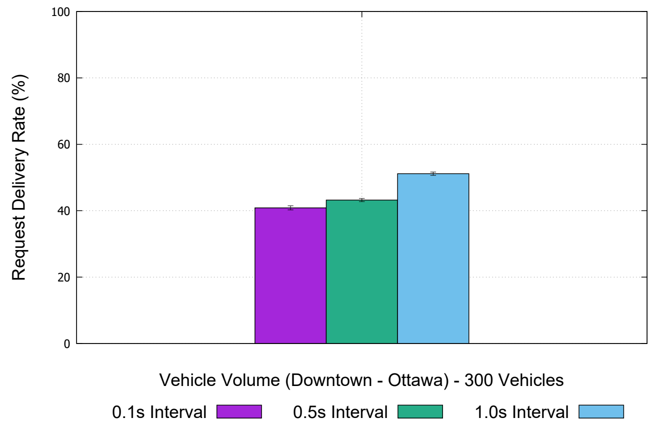

Figure 5.9: Effects of Beacon Frequency in Content Delivery Rates - 300 Vehicle Scenario

network resources decreased delivery rates, with reduction in request fulfillment as high as 10%. As experiments were focused on beacon frequency, the impact of increased beaconing is comparatively smaller to that of increased probability of multimedia transfers, for example. However, it is clear that beacons can negatively affect network operation, causing increased collision and bandwidth wastage. Regardless of comparatively reduced sizes of beacon messages, collision is Moreover, the synchronized intervals in which nodes switch between SCHs and CCH cause beacons to block potential request and data packets from being transmitted. The negative impact of beacons in motivates the maintenance of beacons in a separate window in the CCH.

5.5.4 Class Prioritization Evaluation

In this subsection, we evaluate the effectiveness of content prioritization as a means of increasing content delivery. In doing so we evaluate the proposed LBD prioritization scheme. We also consider the pNDN delay-based prioritization scheme during the evaluation process, with both solutions being contrasted with a baseline defined by the lack of prioritization [90]. The baseline evaluated is based on a standard NDN communication model. In accordance with the prefix system defined we consider content from all three global categories to be present. To focus the evaluation process on the benefits of prioritization of content rather than on the effects of individual characteristics of services, we consider that content from all three top-level categories has the same basic properties. This way, we configured all content objects to be of $100kB$ in size, with each of the three basic categories having fifty thousand (50000) individual objects. Within each content category, objects were distributed according to a Zipf-based popularity ranking and had associated request probability, with the Zipf exponent being set to $\rho = 1$. Content distribution implementation is similar to the previous evaluation approximately 10% of nodes containing any given object in their local caches during system start. The proportion of requests from all three categories was made so that all three categories had the same probability of being selected. Again, this was done to focus the evaluation process in the effect of the prioritization schemes evaluated. Vehicle distribution and density was configured similarly to the previous experiment, with the inclusion of a scenario with 600 vehicles to analyze algorithm efficiency with even higher competition for resources. After a preliminary experiment, the sliding window value to $\lambda = 10$ and packet drop probability to $\mathcal{P}_{\%} = 50\%$.

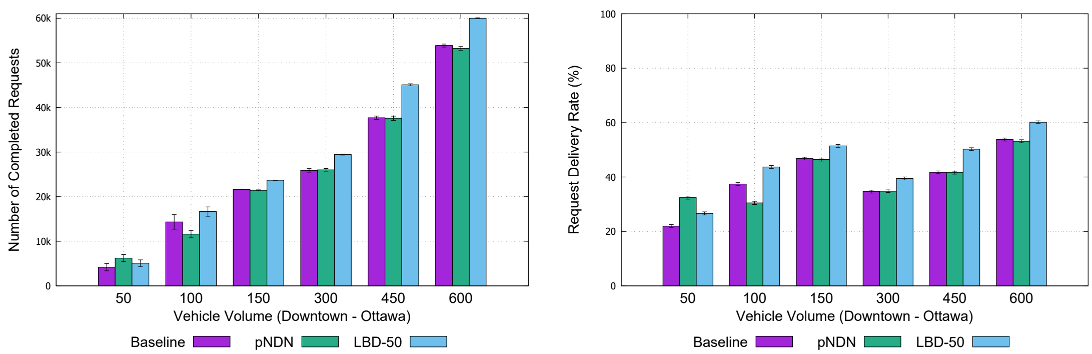

(a) Total Transmitted Data Packets

(b) Delivery rates

Figure 5.10: Total Transmitted Data Packets for the three compared policies in varying traffic density scenarios

Evaluation of the simulation results, denote the benefits of adopting prioritization policies with varying degrees of efficiency. pNDN is proposed as delay-based prioritization schemes and as proposed can lead to increased content delivery rates. However, as denoted in Figure 5.10, this is only observable for low-density scenarios and the effectiveness of the solution is also limited in terms of increase. pNDN allows for approximately 5% in the 50-vehicle scenario, with negligible difference from the baseline no-priority version in all other scenarios. In contrast, LBD promotes increasing gains in terms of delivery rate for scenarios with higher vehicle densities. For the 450-vehicle scenario, LBD provides an increase of approximately 21% over the alternative approaches. Similarly, LBD leads to increases of 13% and 9% for sparser 300-vehicle and 150-vehicle scenarios respectively. The increased gain observed is attributed to its capacity to restrict excessive competition for network resources, as service-deprioritization is only interrupted once nodes perceive that network load has been reduced locally.

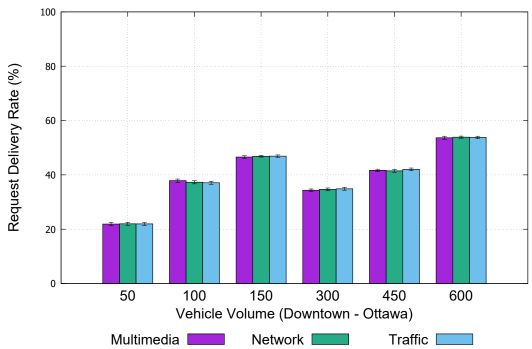

(a) Delivery rates for the Baseline Evaluation

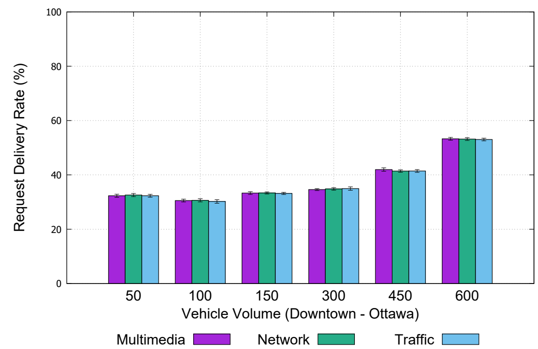

(b) Delivery rates for pNDN : Delay-based Policy

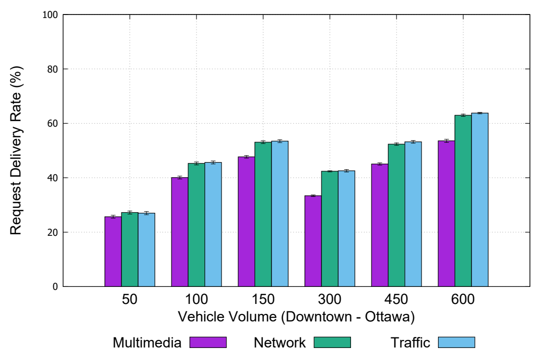

(c) Delivery rates for the LBD : Load-Dependent Policy

Figure 5.11: Delivery Rates for the three different policies on varying network density configurations

The benefits of LBD are further denoted in Figure 5.11, which displays separate graphs with the performance of each policy evaluated. As shown in Figures 5.11a and 5.11b, pNDN has negligible effect in content delivery. While pNDN does not lead to improvements when compared with the baseline scenario, LBD increases delivery of both service categories. As low-priority service restriction is applied whenever load state perception is at the *medium-load* state, it is more likely to be affected than medium-priority services. However, it can be observed in Figure 5.11c that there is a marginal difference

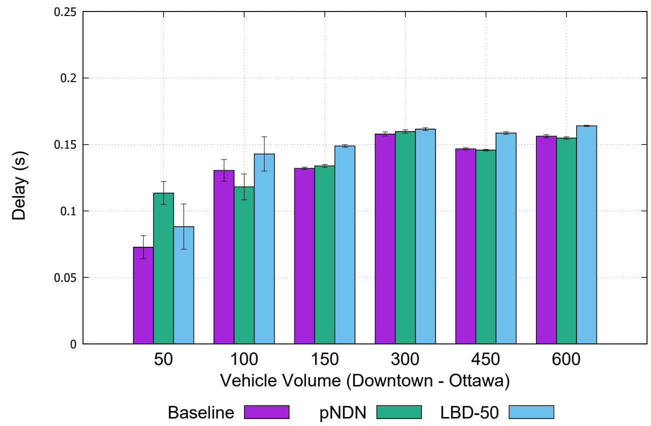

Figure 5.12: Average Latency of Completed Requests for all three tests

between the quality of service provided to medium and high priority services for both the baseline and pNDN approaches. This is a consequence of the smaller restriction-set applied to medium-priority services, which are only partially affected in high-load scenarios. The difference increases along with vehicle density and is predicted to increase as competition for network resources increases, either via increased service request rates or higher network densities. Notably, low-priority service delivery is also increased in LBD. This is a consequence of the "common good" benefit provided by the service restriction policy implemented by LBD. Essentially, by partially restricting multimedia service access, whatever requests are forwarded and fulfilled have higher chance of successfully being fulfilled to the original requester. Overall, LBD achieves the desired goal of increasing both overall quality of service and priority-service efficiency. Lastly, content delivery latency is presented in Figure 5.12. Consistent with its proposal, the application of the pNDN policy leads to decreased latency in the most sparse network configuration, while little change to no difference is observed all other configurations. LBD, however, consistently displays a marginally increased latency, attributed to the increased network use. Notably, the increase is minimal, not being detrimental to solution itself but merely a natural consequence of the benefits offered by it. In fact, given the discussion on the need for service-management and enabled by SEVeN as well as well as increased delivery denoted by application of LBD, the need for further research in service categorization and service prioritization schemes.

5.6 Closing Remarks

In this Chapter, we introduced a novel network model for ICN networks named SEVeN, designed as a service-centric service exchange architecture. We theorized the potential benefits of adopting an information-exchange economy as a mean of controlling information dissemination and enabling service exchange. This concept fomented the SEVeN system architecture, which included the novel service sublayer. SEVeN enables efficient delivery of content model while ensuring that priority services can maintain functionality in face of high network load. It also promotes content sharing as a mean of coordinated access to the limited network services. Our simulations demonstrated how easily network resources are exhausted, furthered evidencing the need for service management. Subsequent evaluation of the LBD service prioritization scheme developed for SEVeN proved effective in reducing average network load and increasing delivery rates in contrast to both the baseline and pNDN technique. In conclusion, we believe that content prioritization schemes are a feasible solution for service management.

Chapter 6

Caching Policies in Information-Centric Networking

This Chapter is structured as follows. Section 6.1 presents an introduction and reasoning behind the focus on Caching in ICN-VANETs. In Chapter 6.2, we present a comprehensive overview of caching in ICN-VANETs, including discussions on open topics and observations on the applicability of ICN caching techniques to the VANET environment. We then discuss the methodology of evaluation we propose to study the impact of caching policies in Section 6.3. In Section 6.4, we perform a series of evaluations on caching policies denoting the effect of parametrized components in content delivery. The Chapter concludes on Section 6.5 with general considerations on caching that guide the works on following chapters.

6.1 Introduction

It is widely known that the greatest challenge of vehicular networks is connectivity. Vehicles generate energy autonomously and are equipped with increasingly ample computing resources, leading to increasing dependence on connectivity. Considering the increasing demand for network services, limited availability of network resources, high level of mobility and intermittent nature of connectivity, it is clear that network resource management is the greatest challenge in VANETs.

As introduced in Chapter 2, by promoting content to first-class citizens and decoupling content from physical addressing, content delivery can be significantly improved in ICN. However, the peculiarities that differentiate VANETs from other ad-hoc networks impose focused research efforts, many of which are currently open topics. One such topic of discussion relates to content caching. ICN caching allows for consistent naming as content travels between source and client with all nodes in the route being capable of identifying the content and consequently acting upon it. This fundamental difference allows for ICN caches to implement functionalities previously impossible such as contextualized in-network caching, decentralized implicit coordination, and chunk level caching. As a result, solutions ranging from safety warning messages, traffic information, advertising to entertainment can benefit from improved content delivery. In essence, most network traffic categories are fit for ICN and applications designed for vehicular environments can further benefit by having the properties of their respective use patterns integrated into ICN-VANET cache solutions. The primary objective of this Chapter is to discuss and evaluate the limitations and opportunities available in the realm of caching as it relates to Information-Centric VANETs, comparing existing works and pinpointing opportunities for improvement.

6.2 Caching in ICN-VANETs

Caching is hailed as one of the prime benefits of named data. ICN networks are expected to exploit prefixing to improve content delivery by employing a diverse number of techniques. Overall Vehicular Network interest in specific classes of contents can be explored to improve delivery of traffic information, for example. In the case of traffic data dissemination, location and temporal requirements need to be evaluated both in terms of data validity, usefulness as well as on where this content is relevant. By structuring naming schemes of traffic information to include the location of the traffic information, the network can enable both directed interest propagation and directed caching. In doing so, content copies can be strategically placed where they are most likely to be requested, in this case in road segments adjacent to the location described by the content. As denoted by this example, knowledge-based systems should take advantage of globally coherent names, user behavior and request patterns to improve delivery rates while reducing network overhead. Such improvements are the expected consequences of in-network caching, which decreases *Average Hop-Count* and *Server Hits* by placing content closer to clients. However, the noisy broadcast environments and highly mobile nature of vehicles pose significant barriers for improvement in VANETs. In the remainder of this section, we discuss the techniques applied to ICN caching and how they are considered in VANETs as well as existing limitations.

6.2.1 Implicit Cache Coordination

Network traffic, cache size and cache lookup time are the most commonly considered parameters in the definition of cache policies. Psaras et al. proposed defining caching for ICN in terms of network activity, measuring the available length of time that can be cached given cache size and network bandwidth and request frequency [72]. Based on these metrics, the authors proposed a probabilistic caching algorithm called *ProbCache*. The authors argue that fair resource allocation should be performed based on resource sharing. However, one of the main limitations in vehicular environments is the temporary nature of network connections. The high level of node mobility in a VANET leads to short-lived connections and consequently to limited path related traffic flow measuring ability. Thus, estimating flow volume in intermittent connectivity environments subject to flow bursts requires an alternate approach. Traditional implicit coordination schemes not dependent on network state can be applied without modification in ICN-VANETs; however, their effectiveness will be limited. Leave Copy Down *(LCD)*, Move Copy Down *(MCD)*, Copy with Probability *(Prob)*, Random Copy One (RCOne) [71, 68].

6.2.2 Explicit Cache Coordination

In an ad-hoc environment, cache coordination is heavily dependent on network stability. When multiple nodes can coordinate cache distribution, optimal cache strategies should distribute content storage to allow for a higher probability of content being available. While this is trivial in fixed topologies, efficient coordination is a very complex task, especially in high mobility networks including VANETs. Centralized caching in VANETs relies on on vehicle-to-infrastructure (V2I) interaction. Mobile nodes interact with Road-side Stations (RSUs) access points or directly to Base Stations (BSs). Cache policies are maintained by the BSs and can be coordinated with RSUs to predictably cache content a given vehicle behavior. Mobility prediction is often cited as an example of infrastructure based caching policy not only being applied to ICN networks but VANETs in general. The potential benefits of centralized cache coordination are dependent on the deployment of centralized control structures. Creation of such structures is extremely costly, and as discussed previously, regardless of the existence of centralized structures, efficient network management is dependent on collaboration between all players involved. While future standards may include additional infrastructure features, they will be capable of operating independently from dedicated vehicular infrastructure services. Notable examples include 5G and WAVE, designed as a general application architecture and an ad-hoc friendly architecture respectively. Hence, the efficacy of centralized caching in ICN-VANETS relies on accurate knowledge of vehicle mobility patterns as well as integration with V2V systems [30].

6.2.3 Chunk-level Caching

Universal naming allows for partial content storage which can further reduce communication overheads and can be incorporated into both implicit and explicit caching policies. Implementing partial caching as part of a predictive caching scheme allows for only the required portions of a file to be placed at any given location. Partial caching can also reduce cache failure probability by taking advantage of client behavior. Chunk-level popularity can also be exploited to store partial copies of content. One such example relates to multimedia content and social media. Video file segments are sequential, and initial segments are more likely to be viewed, allowing for caches to consider popularity when choosing what portions should be cached. However, by caching object segments, cache lookup can be significantly increased. To avoid increasing storage and lookup overheads, most solutions maintain a single reference all chunks belonging to the same object. Chunk storage efficiency depends heavily on content granularity and correlation. The access frequency of individual chunks will differ and similarly to objects, content chunks follow will follow some distribution, very commonly it being the Zipf distribution [105]. It should be noted that file popularity cannot be extended to the chunk level as they follow independent distributions. Hence, implementing chunk-level caching requires additional knowledge, harder to maintain in systems such as ICN-VANETs and whose implementation poses additional challenges.

6.2.4 Predictive Caching

Predictive caching takes advantage of network knowledge to improve content delivery. The patterns considered are generally based on content popularity, spatio-temporal request correlations, and mobility patterns. In real-world applications, content interest is mutable and Poisson processes rarely describe the distribution of requests. ICN-VANET networks exhibit particularly dynamic behavior, being extremely volatile as a result of highly mobile topology, the ubiquity of in-network caching and lack of coordination already prevalent in ICN environments [29]. As vehicles have a fragmented knowledge of the network, predictive caching in V2V environments requires cooperation and knowledge sharing, ultimately being dependent on the connection time (also known as *sojourn* *time*) vehicles have to their neighbors. This way, predictive caching is targeted towards infrastructures or vehicle groups.

6.2.5 Mobility-aware Caching

Mobility-aware caching is a specialized predictive caching mechanism for both V2V and V2I systems used to store content close to vehicles based on expected mobility patterns. It allows infrastructure to place caches at the RSUs closer to nodes regardless of their movement. Average hop distance and network load should consequently be decreased proportionately to the availability of RSUs and accuracy of movement prediction. Chunk-level caching can be incorporated as part of mobility-aware schemes to reduce The two greatest challenges in the implementation of mobility-aware caching policies are the privacy concerns and the need to deploy specialized infrastructure. Regarding the privacy aspect, it has been shown that vehicles can be tracked with even with limited information and pseudo-randomized user identifiers [106]. Given issues such as impersonation, privacy concerns can potentially hinder the viability of implementing mobility aware infrastructure solutions [107]. In V2V environments mobility-aware caching is significantly limited, and mobility patterns are used as knowledge in the definition of cache coordination policies.

6.2.6 Adaptive Caching

Global mobility patterns, as well as other periodic behavior, can be used to improve caching at RSUs. Recurring events such as rush hour can be predicted, allowing for contextually popular content such as traffic, weather or advertising to be cached. Notably, implementing adaptive cache policies is dependent on a deep understanding of group behavior and spatio-temporal correlations between requests.

6.2.7 Spatio-Temporal Correlations

By exploiting location or time-related correlations, content can be cached where or when they are more likely to be requested. Unlike other application specific contexts, spatio-temporal properties are universal even if not used by requests and can be incorporated more broadly into solutions. A common example of location dependent content is traffic information, relevant to the specific location that it describes and whose freshness varies according to the traffic pattern. In high-traffic scenarios, it is likely that a location will become a hotspot for traffic requests. Spatio-temporal data can also be used as freshness factors and be parameters considered by cache eviction policies. However, it should be noted that as cache in ICN needs to operate at line speed, it is suggested that cache policies be kept minimal. Additional heuristics can be implemented to take advantage of spatio-temporal correlations between requests but those are dependent on application context.

6.2.8 Social Caching

Structuring a network as subsets of groups has the potential to reduce overall network overhead and allow social groups to coordinate content caching. The division of a network into small groups is heavily dependent on inter-user contact time but allows for distributed explicit cache coordination. Social groups enable the observation of recurring behavior within a V2V context. Socially generated knowledge allows for additional routing policies as well as prediction of group behavior to be explored. However, the efficacy of social policies is dependent on vehicle contact time and is targeted for more cohesive vehicular contexts such as traffic jams and highway routes.

The greatest limitations to the implementation of caching policies in VANETs are the peculiarities that separate Vehicular Networks from general MANET. The high level of mobility, the current lack of centralized structures, limited connectivity, constant topology changes and current stage of research are fertile ground for the definition of new solutions. To understand the properties of ICN-VANET environments, we proceed with an evaluation of currently existing solutions and aim to evaluate the effects the selected policies on VANET environments.

6.3 Policy Evaluation Methodology

To evaluate the effects of the different components of caching in ICN-VANETs, we study the effects of each parameter selected in simulated network scenarios. Given the nature of the experiments and limited availability of real-world traces, we resort to simulation-based evaluations using well-established simulation tools. We implemented and evaluated multiple parameters including varied library sizes, content sizes, popularity distributions, replacement policies, cache coordination policies and network topologies, as well as zipf characterization exponents though we do not display certain results due to space limitations. For a portion of the work, we briefly mention results obtained and their relevance in the evaluation. The ICN-VANETs In-network caching models were incorporated into the Veins simulation framework (*veins.car2x.org*), OMNeT++ network simulator (*omnetpp.org*), and SUMO road traffic simulator (*sumo.dlr.de*). To test the various caching related techniques considered, additional functionalities were added to the simulation environment. Using the configured setup, we gathered results from thousands of simulations representing multiple iterations of each scenario configuration.

6.3.1 Information-Centric Networking

In ICN nodes consume and provide content based on globally known named data objects. Content requests are generated by clients and are labeled as interest packets, which are checked locally and forwarded should the content not be available locally. While there are many proposed approaches to how content delivery should be implemented, including DONA [20] and NetInf [108], we adhere to the most well known and most popular solution, CCN [19]. In a CCN network, interest requests are forwarded throughout the network and at every step are checked against the Content Store (CS) cache, the Forwarding Information Table (FIB) and the Pending Interest Table (PIT). Interest packets are routed throughout the network towards potential content routers, which reply with Data packets. Data packets follow reverse routes towards the client and at every step are checked against the local PIT which can use the same data object to fulfill multiple pending requests. The unified naming scheme and flexible interest/data packet types allow for significant improvements in caching a data packet will always be able to fulfill interests for a matching identifier. During the remainder of this text, when referring to the ICN approach consider, we refer to the CCN design model as well as to elements of its specific terminology.

6.3.2 Parametrization

To evaluate caching from a wide ICN-VANET context, we chose to study the effects of existing in-network caching policies and their effects on the network. This allows us to evaluate cache behavior within the constraints of the space available. In doing so, we chose to evaluate the influence of *content popularity*, *cache sizes*, *request frequency*, *content replacement policy* and *cache coordination policy*. Considering popularity is fundamental as it dictates the frequency in which individual objects are requested one of the fundamental motivations behind studying caching. In one hand, purely randomly requests reduce the efficacy of caching policies, while on the other should a limited number of requests be overly dominant caching policies can become redundant. Evaluations considered the following values for the Zipf exponent: $\{1, 2, 3, 4\}$. The cache size plays a fundamental role in the efficacy of the caching policy. While a cache size should me limited to avoid increased lookup times and maintain line-speed operation, caches should be able to maintain a minimum set of items that justify its existence. Individual content objects sizes were set to *10kb* and caches were configured to hold the following number of items: $\{5, 10, 15, 25, 50, 75, 100, 125\}$. The cache size was selected to be proportionate to the number of unique items available, which was set to 50.000. Three replacement policies were evaluated: *Random replacement*, *Least Frequently Used* and *Least Recently Used*. Network coordination was evaluated considering the coordination policies discussed in Section 6.2.1. Additionally, requests frequency was set to an average of 0.75 seconds and simulation time was set to 1500 seconds, ensuring that the network has enough time to reach a ready state. Lastly, it is important to note that the parameter values were selected based on related work in ICN and on experimentation, as there is no similar work available for VANET environments.

It should be noted that given the current stage of research on ICN networks, limitations such as the lack of a standard architecture, additional constraints imposed by the vehicular environment as well as limitations in naming schemes, restrict the evaluation of chunk-level caching. As naming is application dependent, a separate study targeted towards naming schemes is suggested. One such work would focus on evaluating cache granularity, aggregation and the effect of chunk-level updates on cache efficiency. In this work our focus is in evaluating the applicability of currently existing solutions, targeting the parameters listed in this section, future efforts will target chunk-level caching.

6.3.3 Network Topology

To evaluate caching we consider two different contexts. The first configuration relates to representations of clusters, where vehicles are tightly connected. These configurations are represented by a "static" network with a stable topology. The grid-like scenario was selected as it offers a way to compare results of the wireless networks to wired networks.Increasing the number of connections in the set would degenerate it into a single-hop network, limiting evaluation of content forwarding and intermediate caching, further denoting the relevance of this configuration. The structure modeled is displayed in Figure 6.1a where both potential symmetrical and asymmetrical content provider locations are displayed. The group structure also represents how groups can behave within

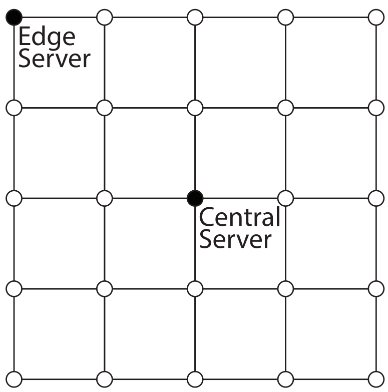

(a) Wireless Grid-like Network

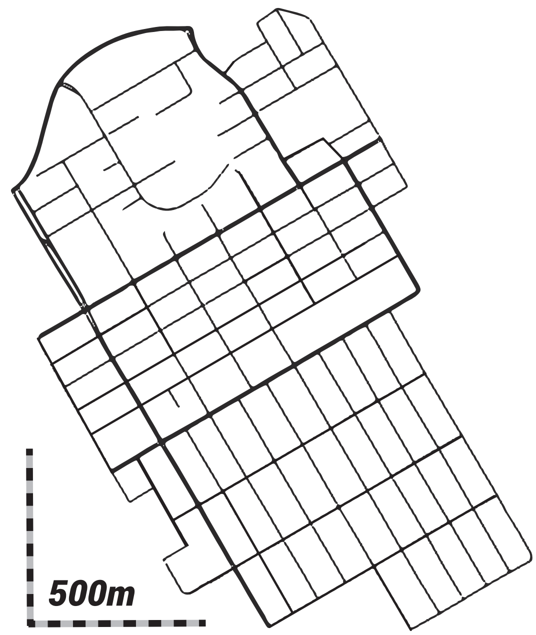

(b) Downtown Area Map

Figure 6.1: Simulation Scenario Configurations

a network. The second scenario considered represents an urban environment, where vehicles move freely independent of one another. The urban environment scenario selected was based on a real-world downtown location used in Chapter 5. The configurations selected for testing allows for evaluation with varied traffic flow densities.

6.4 Simulation Evaluation of Caching Policies

This section presents the evaluation of the simulation studies. First, we briefly investigate the effects of content popularity on caching. We then evaluate both scenarios proposed in Section 6.3.3. We evaluate the effects of the parameters defined in Section 6.3.2 in terms of average hop distance, server hit saves, delivery rate, and communication latency.

Prior to a wide-scale evaluation of the scenarios proposed, we briefly evaluate the effects of content popularity. Simulation results denoted an extremely asymmetrical distribution of content for Zipf exponent values above 1. The lack of variety in content request, a consequence of the properties of power-law distributions, lead to a small fraction of items rapidly occupying caches and negates potential benefits of caching policies.

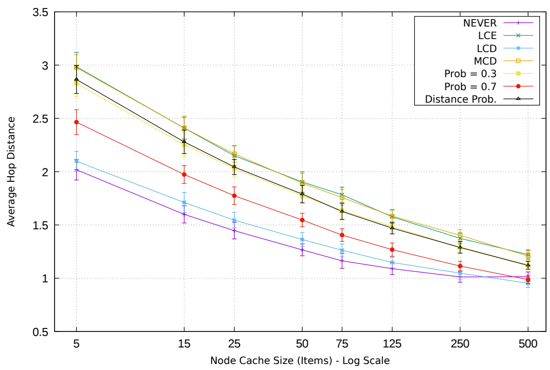

(a) Edge Server - Average Distance

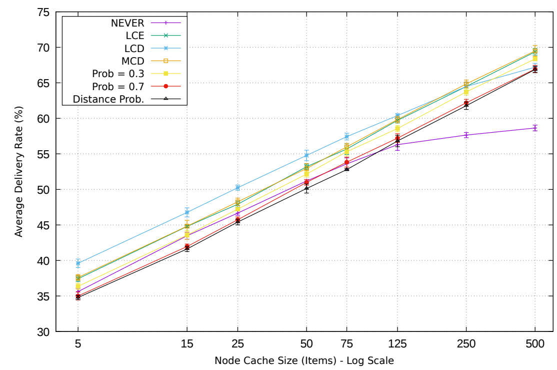

(b) Edge Server - Average Delivery Rate

Figure 6.2: Simulation Results for the *asymmetrical* Grid Scenario

Hence, for the remainder of the work, we consider the Zipf exponent value to be 1, a trait shared with other works that investigate caching. Similar considerations are made to the cache replacement policies. Initial evaluations also demonstrated a small increase in success rate for frequency based replacement over purely random cache replacement, with improvement being more pronounced for Zipf$=1$. These results supported the selection of LRU as the replacement policy of choice for the remainder of the simulation evaluation process.

6.4.1 Grid-like Network Scenario

Results for the asymmetrical configuration are displayed in Figure 6.2. Figure 6.2a presents the average distance in hops for all completed transfers. The results displayed denote significant advantage in adopting protocols with less likelihood of performing cache inclusions. This is the result of the relatively high connectivity degree of the network in conjunction with the asymmetrical position of the content server. In this scenario, policies that are eager in performing cache inclusion (and subsequently removals) fail to maintain copies of requested items locally, increasing average hop-count. Notably, the *Leave Copy Down* policy has results comparable to the *Never Cache* policy, a consequence of its slow cache dissemination speed and medium cache redundancy. The benefits of the LCD technique are corroborated by Figure 6.2b, which displays overall delivery rates. The results denote that LCD achieves a 5% increase in delivery rate when compared to the *Distant Probabilistic* (DProb) and a 25% decrease in average latency. *DProb* tends to place content closer to clients, yet it performed worse than the randomly

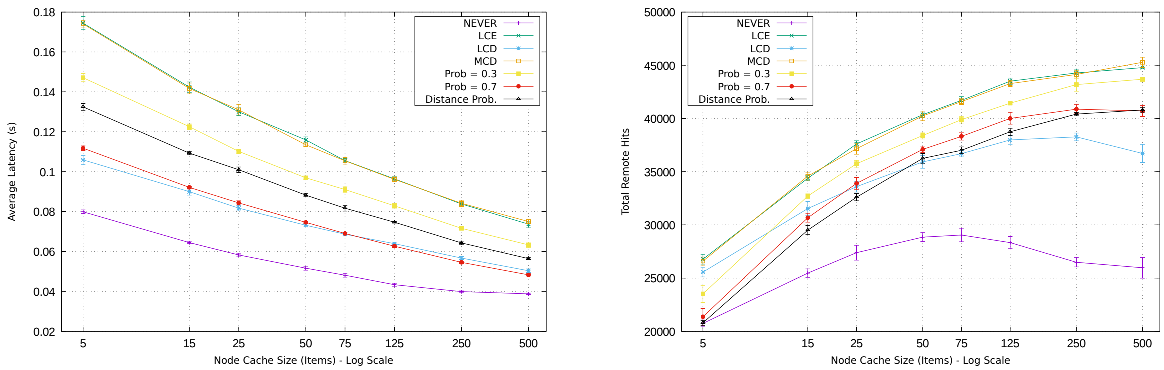

(a) Central Server - Average Latency (b) Central Server - Total Remote Hits

Figure 6.3: Simulation Results for the *symmetrical* Grid Scenario

probabilistic approaches with higher chances of adding an item to caches. While *DProb* has a smaller average probability of caching content than the probabilistic approach with 70% cache probability it still yields lower results specifically because contents caching is necessarily asymmetrical. This further denotes the lack of spatial correlation in request locations. By reducing propagation speed of cached copies, the system reduces the stress in the links close to the server, increasing the total delivery rate at the cost of average delivery time. While this implies an increased network cache redundancy, the reduction in total network messages and increased delivery rate offset the potentially reduced cache variety.

Results for the symmetrical configuration denote similar trends to the asymmetrical case, and are displayed in Figure 6.3. In this configuration, the network load is evenly distributed and average distance to the content server is halved, leading to increased server request probability. Given the average distance of 2.5 hops to the content server, policies that place content close to clients fail as well as the LCE policy fail to use cache efficiency and decrease cache variety. *Total server hit* was reduced as a function of cache sizes and due to the centrality of the server node. Average *hop-count* is reduced by an average of 0.5 hop, a poor result given that the average distance to the server is reduced 1.5 hop between both scenarios. The delivery rates follow a similar distribution to the ones in the asymmetrical scenario. The mismatch between reduction in hop distance as well as latency pinpoint the benefits of adopting slow propagation cache solutions and that solutions that place content closer to previous requesters are not necessarily beneficial. Lastly, Figure 6.3b denotes the effect of cache size in node capacity to enable intermediate caching. The effect of popularity in power-law distributions is exceptionally

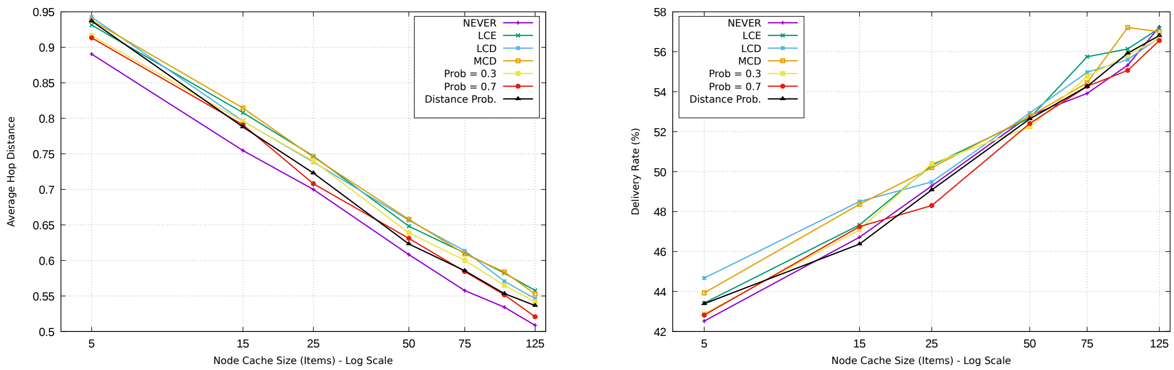

(a) Downtown - Average Distance

(b) Downtown - Average Delivery Rate

Figure 6.4: Simulation Results Downtown Scenario

visible for *Never Cache* policy, where increased cache size do not lead linear growth in intermediate cache hits. In this case, as vehicles are only capable of serving requests for content they themselves have requested in the past, expanding cache size increase cache homogeneity, increasing local cache probability and, therefore, reducing intermediate caching as denoted by the reduction in remote hits for caches capable of storing more than 75 objects.

6.4.2 Downtown Area Scenario

Based on the results obtained we proceed to the evaluation of a real-world scenario. For this scenario configuration, 10% of the nodes were configured as content providers. Figure 6.4 displays the results obtained for a vehicle density of 25 *vehicles per* km^2. Results pinpoint a smaller difference in efficiency between solutions. This is the result of the volatility in vehicular environments. As vehicles move independently, the network connectivity structure is constantly changing and the negative impacts of placing content close to the original requester are mitigated. However, the benefits of slow propagation are still visible. Another relevant feature observed is the stability of delivery rates. Regardless of network structure, all scenarios display similar delivery rates. Despite a reduction in overall network connectivity, the downtown scenario preserves network connectivity allowing for comparable content delivery. Notably, vehicle mobility leads to variations in proximity to content servers, affecting average server hit numbers. These results do not affect global simulation results, as decreases in server hit are offset by increase remote hits. Additional simulations with increased vehicle density produced similar results.

6.5 Concluding Remarks

In this Chapter, we discussed the intricacies of caching in ICN-VANETs. We pinpointed potential techniques that can be incorporated and the challenges that need to be overcome. We evaluated a series of properties with the goal of presenting the limitations of existing solutions. We demonstrated that in the context of VANETs, considering mobility is paramount and that eagerness in caching content does not necessarily lead to improvements in content delivery. We also denoted the importance of cache size in comparison to other properties.

Chapter 7

Utility-Gradient Implicit Cache Coordination

This Chapter is structured as follows. In Section 7.1 we contextualize the work presented in this Chapter. Following, Section 7.2 presents an overview of caching in ICN-VANETS. Existing solutions and the conceptualization for the model proposed are made in Section 7.3. Following, in Section 7.4, we perform a series of analysis on the proposed caching mechanism, comparing it to existing solutions and validating our model via simulation. The document concludes with considerations on potential improvements and future works, presented in Section 7.5.

7.1 Introduction

Increasing connectivity and demand for services put significant strain on vehicular network services. The development of an Internet of Vehicles (IoV) and competition for mobile telecommunication services further burden the connectivity capabilities of vehicles. The rise of service-centric VANETs and internal content production-consumption cycle motivate the introduction of new network paradigms where content becomes a central figure and where vehicles cooperate to decrease the wireless network load and dependence on infrastructure services. The fast growth of user-generated content and multimedia services further denote the need for efficient network access. Physical addressing is not designed for highly mobile networks, and increases in access contention further strain its operation. User mobility and the IoV motivate the adoption of Information-Centric Networking (ICN). Unlike physical addressing, ICN is tailored for the VANET environment, with named data being a perfect fit for an ever-increasing vehicular application space. Caching is one are where universally understood content names allows for significant improvements. As vehicles gain awareness over the content broadcasted within the network, they can operate on it to improve content delivery and reduce network load. To manage the production and consumption of content within a VANET, ICN takes full advantage of content naming in its cache mechanisms. Named data allows for significant improvements in intermediate caching in mobile networks, which is a fundamental requirement for VANET efficiency [33]. Content popularity estimation is one such way for vehicles to improve their cache efficiency. By predicting which items are most likely to be requested, nodes can improve cache hit. Mobility, however, dampens network capacity to coordinate caches, forcing vehicles to maintain their caches with limited information. Implicit cache coordination via information exchange embedded as part of content delivery mechanisms allows for increased awareness and subsequent refinement of cache policies.

7.2 Caching in ICN VANETs

A significant portion of ICN literature tackles communication from an infrastructure to mobile device standpoint. Recent studies on Infrastructure-based mobile ICN have considered the properties of ICN to implement sophisticated algorithms and protocols. Aspects such as content popularity, vehicle mobility, spatio-temporal correlations and social patterns to design their solutions. The PCP cache insertion policy for NDN networks is one such example, proposing a handover mechanism for content delivery and taking content popularity as a factor in caching [109]. Similarly, the Mobility/Popularity-based Caching Strategy (MPCS) proposes a network fallback mechanism to assist in content retrieval [110]. However, while the authors consider mobility and popularity in their solutions, the proposed schemes are infrastructure dependent. Ad-hoc ICN, however, is faced with additional limitations that require infrastructure-centric solutions to be adapted or for novel ad-hoc-centric alternatives to be designed. Limited and intermittent connectivity, partial network perception and partial knowledge are examples of challenges that require work. Content chunking, however, is a subtopic where solutions apply to most ICN systems without the need for significant adaptation. By segmenting content transmission and storage while maintaining context, nodes can partially cache objects. WAVE and ICC, for example, are two solutions that target ad-hoc networks and implement popularity-aware chunk-level policies focusing on access count and social distance interest prediction respectively [111, 112]. While both solutions, as well as additional related work, consider ad-hoc networking, they do so at the cost of network structure flexibility. Such solutions impose limitations to aspects such as network size, node mobility, content sizes and request frequency distribution. This way, while content chunking is an important feature and should ultimately be part of any ICN cache system, it is not solely capable of solving cache in information-centric networks. State of the art on the topic denotes limitations in cache coordination and insertion policies targeted towards ad-hoc networks and, more specifically, vehicular networks. Surveys on the subject of Mobile ICN evidences the limitations in the literature, with a limited number of solutions being applicable in vehicular environment contexts [113]. This research vacuum provides the opportunity for works focusing cache insertion solutions designed for ICN VANETS.

Unlike cache eviction policies, which rely on local cache state, insertion policies operate with network context. By explicitly exchanging information or by implicitly embedding knowledge in request responses cache insertion can greatly improve cache efficiency at a network level. Cache insertion policies based on explicit coordination allow for greater knowledge cohesion across the network regardless of network state at the cost of increased overhead. The increased awareness allows for improved cache coordination, reducing cache redundancy. Exploring social relationships, exemplified by the related work on chunk-level cache, is one of the many approaches to explicit cache insertion coordination. Despite its potential benefits, explicit coordination relies on connection persistence with social policies also depending on request pattern correlation between nodes. Implicit coordination, on the contrary, does not rely on additional communication and has a universal appeal. The potential benefits of implicit coordination might be limited when compared to those of explicit coordination, but it's fewer drawbacks are certainly preferable. Within the context of implicit cache insertion coordination policies, existing solutions are significantly limited in their exploitation of network knowledge, as discussed in Chapter 6.

Content can be obtained from sources either internal to the VANET or externally from base stations or other infrastructure devices. However, it is suggested that resorting to infrastructure should be considered a *last-resort measure*, given potential the costs associated with offloading. Hence, it is preferable that content is obtained from within the vehicular network whenever possible. Moreover, given that VANET both network resources and cache sizes are limited, cache policies should be designed to minimize network usage while maximizing cache hit rates. Cache Policies should reduce the need for network access. Existing cache solutions designed for ad-hoc networks implement systems where cache placement does not take into consideration two factors fundamental for mobile networks: content popularity and cache distribution. Amongst existing solutions, Cache Everything Everywhere (CEE), Never Cache, Leave Copy Down (MCD), and Probabilistic Cache (DistProb) are the most notable [71, 68]. Note that existing solutions either make either none or limited consideration to properties such as content popularity.

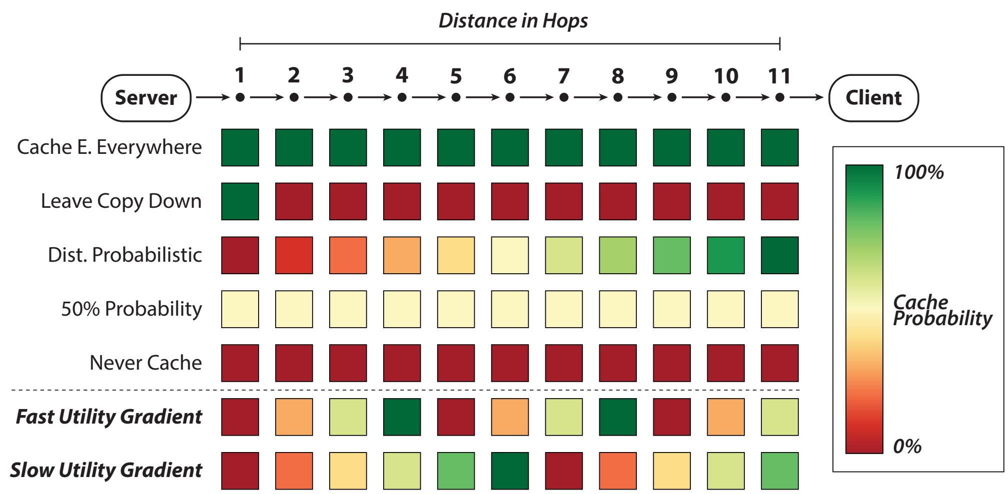

Figure 7.1: Cache Placement Probability for Multiple Cache Policies

Fig. 7.1 displays an example of how likely a content object is to be cached at subsequent intermediate nodes during content delivery. In this example, each box displays the likelihood of the sample object being added to the cache of intermediate nodes with varying distance in hops from the original sender. From the example, we can observe the behavior of different policies. The CEE policy forces cache replacements for all requests, leading to high cache replacement rate with potentially high cache redundancy depending on the eviction policy deployed. Constant yet blind cache replacements limit application in scenarios with skewed item request frequencies. Conversely, the Never Cache policy strictly prohibits intermediate caching, leading nodes to only store objects they themselves have requested. The LCD policy only permits cache replacement at the first intermediate node from the server, reducing cache redundancy for low-frequency items while increasing it for frequently requested objects. This behavior mimics a popularity filter, behaving similarly to the WAVE chunk caching algorithm but with linear propagation speed. LCD, however, does not explicitly take popularity into account and only increase replication for items with request frequency above a $1/cacheSize$ threshold. The DistProb policy makes cache copies likelier to be placed closer to clients, increasing redundancy asymmetrically. The Distance-probabilistic cache policy considers the distance from sender to the receiver but makes no consideration regarding content popularity or cache duplication, causing copies to be placed asymmetrically closer to the client. The lack of a tightly correlated location-dependency pattern makes this policy ill-fitted for most scenarios. From the limitations pinpointed, we propose the definition of a policy that distributes cache placement throughout return paths to avoid localized cache duplication while allowing for increase cache variety. The work proposed is based on a *Utility Gradient*, discussed next. In the following section, we present the proposed model, its features and how it is designed to take advantage of network knowledge in a way that existing works have not.

7.3 Proposed Solution

To improve cache coordination efficiency, we tackle the cache management problem by treating it in terms of increasing the probability that requested items are cached and minimizing the average hop distance to content, should it be cached within the network. Continuing the example presented in Fig. 7.1, we propose a utility gradient function, designed to inform nodes performing intermediate caching on the potential a content object has to increase cache hit rate. In the example displayed, two possible versions of the proposed model with different λ distance factor are listed. One example makes nodes more likely to suggest cache insertion while the other makes it less probable for them to perform the same action. The two variants, denoted as "fast" and "slow" gradients, perform the same basic evaluation of the usefulness of an item to the network compared to items currently available in their caches for different λ values. By having content providers share the request frequency for requested items and having nodes in the return path evaluate this information, nodes can decide to cache items with a higher utility to the network, compromising a balance between increased cache variability and caching the objects with greatest probabilities of being requested. The distance from the content provider (server) to the node where last cache insertion operation was performed for the current request is also considered in the evaluation to avoid having multiple copies be placed next to one another. The purpose of this approach is to distribute content objects throughout the network more evenly while taking popularity and cache distance into consideration in its operation. For the remainder of this chapter, when referrals made to cache distance relate to the distance between the node where the cache insertion policy is being evaluated and the last node where the current item was cached, given in hops.

The distance is computed based on the return-path and is updated at every intermediate node that participates in the request response path. Regarding the example provided, the goal of the policy proposed is to spatially distribute copies with cache distance being proportional to item popularity as well as additional parameters such as the λ distance factor. The communication model adopted for our cache solution is based on the NDN model, which operates based on interest propagation [9]. The algorithm proposed and its utility function is discussed next.

7.3.1 Utility Function

To formulate our utility function, we denote the following elements:

- $\mathcal{N} : n \in \mathcal{N} = \{1, 2, \cdots, N\}$ denotes the group of nodes (vehicles) in the network with a group size of N.
- $\mathcal{C} : c \in \mathcal{C} = \{1, 2, \cdots, C\}$ denotes the list of content objects of size C available in the network.
- Z_i : Cache size of node $i, i \in N$. For the sake of simplicity, we consider that all nodes have the same caching capacity Z, measured in the maximum number of equally-sized objects that can be cached, or $Z = Z_i = Z_j; i, j \in N, i \neq j$.
- P_i : popularity of content the i-th object in the popularity rank sequence whose most popular item is ranked as #1 and presented in terms of its probability mass in terms, given in terms of the items request frequency.

For a network with a single node, the caching decision is rather trivial, with a node caching the items most likely to be requested. While in the real-world scenarios, nodes would have to rely on counting requests to infer item frequency, a more formal analytical evaluation can be made considering request frequency distributions. In the work in this Chapter, we consider object request frequency to be described by a Zipf-like distribution [105], a power-law distribution frequently considered in the literature, and given as:

$$PMF(x; \rho, C) = \frac{x^{-\rho}}{\sum\limits_{i=1}^{C} i^{-\rho}} \tag{7.1}$$

where x is the P_x-th ranked item, and ρ is the exponent that characterizes the Zipf-like distribution, it is optimal to cache the highest ranked items. By having the only available node cache the most popular content objects available in the network, local cache hit probability is maximized. This increase in hit probability occurs regardless of the node being capable of caching all existing objects or only a subset, considerations are made to a single cache. In the case where there are two or more nodes in the network, caching decision becomes complex. Policies must consider if the number of existing objects C is greater than the combined caching capacity of the existing nodes and if the nodes are capable of caching all available objects as a group. If it is not possible for the network to cache all existing items, the only viable solution is the definition of schemes to maximize cache hit probability and minimize network overhead. Otherwise, each node is individually capable of caching all existing objects ($Z_i > C, i \in N$), the case degenerates to the first example, with each being independently capable of managing its cache. However, if nodes are individually not capable of caching the content coordination is necessary ($Z < C; N \cdot Z \geq C$), to maximize the probability of a cache hit, content has to be distributed to manage cache redundancy for as many objects to be cached. For each node to maximize their local probability cache success rate, content needs to be distributed in a coordinated manner. Content distribution should attempt to maximize the likelihood of content being locally available while ensuring that all items are cached somewhere in the network. In an optimal solution, cache redundancy should only occur if the total network cache space exceeds the number of cacheable objects C. For $N \cdot Z = C$, cache distribution is such that the difference between the cumulative probability of items cached at each node is minimized. Yet, the number of combinations in which C objects can be divided amongst N nodes is factorial. Hence, finding the combination that minimizes the difference in cumulative probability is a non-polynomial problem. Similarly, maintaining a global coordination scheme is not feasible given the overhead incurred and likely connectivity limitation vehicles will face. This way, ultimately for both (i) and (ii) the fundamental question is where to place items to benefit both nodes individually as well as the network as a whole.

Individually, nodes gain by caching the most popular items locally while as a community by increasing cache diversity. To compromise between these contrasting priorities, we propose that items be cached with frequency relative to their popularity. The proposed solution, dubbed Utility Gradient Cache (UG-Cache), is a passive coordination scheme that takes advantage of network knowledge to improve cache efficiency similarly to what is done in active schemes but without the need for explicit knowledge exchange. This is possible given the inclusion of additional meta-data in responses to content requests, describing basic characteristics of the objects' access pattern. By providing the request count and relative frequency of requests for the desired object, nodes in the return path can roughly estimate the popularity of the content object and determine its potential for inclusion in their caches. To avoid having subsequent nodes coming to the same conclusion and deciding to cache the item we consider the distance between the last cache operation as an additional parameter included in the return path and updated at every intermediate node independently from its decision. The resulting item utility function is as follows:

$$\frac{\left\lfloor -\log_2 (d+1) \cdot r \cdot \lambda \right\rfloor}{N \cdot \sum_{i=1}^{Z} P_i}, \tag{7.2}$$

where d is the distance in hops between the current node and the last node at which the object was cached in the reverse path, r is the Request count for the requested item at its source location, and λ is the distance impact factor. Equation 7.2 is then evaluated as part of the Insertion Decision Algorithm. At every intermediate node, if application of 7.2 yields a value equal or greater than 1, or if the cache is not full, the item should be added to the cache, triggering a replacement. The resulting inequality is dependent on simple information and has a small computational fingerprint, being fit for ICNs, which feature increased cache participation and require line-speed operation for caches. During cache insertion, item request frequency is also stored to be used in the utility equation in future requests, and subsequent cache hits increment the lookup count. Notably, we do not enforce item selection for replacement as part of the cache insertion policy as this is outside the scope of this Chapter and as it merits additional research, as listed in Section 9.2. Additionally, to allow for a ampler model proposal, we consider a variant of the proposed solution where the item with the lowest request frequency (highest rank) cached is used in the equation instead of the local cache popularity sum for completeness. The Equation for the minimum frequency variant is given as:

$$\frac{\left\lfloor -\log_2 (d+1) \cdot r \cdot \lambda \right\rfloor}{N \cdot \left(\min_{\forall i \in Z} P_i \cdot Z\right)}, \tag{7.3}$$

The minimum frequency variant, which we refer to as the Minimum Gradient Cache (MG-Cache) is presented as an alternative, designed to further increase potential for cache replacement given the context of limited knowledge but more dependent on current cache state rather than globally as a result from changes to λ. MG-Cache behaves similarly to UG-Cache with the caveat of it being dependent solely on the properties of the least popular content object cached locally. For both variations of the proposed algorithm, we consider the *Least Frequently Used* replacement policy as being deployed. Maintenance of request counts is implemented using temporized sliding-windows for individual content objects. Given the limited cache size, the additional space required for frequency estimation is not significant. Moreover, computing request frequency and maintaining sliding window timers can both be made asynchronously and independently from other functions of the communication process. Next, we perform a basic numerical evaluation of the proposed models prior to simulation evaluation.

7.3.2 Cache Insertion Analysis

From the proposed utility function, one can infer that cache eligibility is tiered depending on distance and item popularity. To evaluate how the equation behaves we can consider the Probability Mass Function of the Zipf-like distribution considered to model item request frequency mentioned previously. Given 7.2 and 7.1, we define a generalized Utility Function, defined as the following inequality:

$$\frac{-\log_2(d+1) \cdot \left({x^{-\rho}}/{\sum_{i=1}^{C} c_i^{-\rho}}\right) \cdot \lambda}{Z \cdot \sum_{j=1}^{Z} \left({j^{-\rho}}/{\sum_{i=1}^{K} z_i^{-\rho}}\right)} \geq 1 \tag{7.4}$$

With the generalized equation, we can evaluate what items fit the function for different values of d, r and λ. In our evaluation, we consider a restriction based on role of ρ in the distribution, having $\rho = 1$. Increases in the exponent lead to greater request probability for higher ranked objects. Consequently, request diversity is reduced and the need for efficient caching becomes less apparent. In doing so, we adhere to the existing literature which forgoes values greater than 1.

The Ranking of items that fit the generalized inequation used in our algorithm is presented in Fig. 7.2. The graph displays connected line-point segments for different cache sizes Z, $\lambda = 1$ and a content library of 2000 objects. For each line, the ranking of the item with the highest rank that fits the inequality is plotted for varying distances from a theoretical content source. From the graphs, it can be observed that increases in cache size have increasingly limited effect in the increase of item rank viability. Relative to cache sizes, items with ranks up to 4× are viable for inclusion. Additionally, for the distance of 1 hop, rank viability values match cache sizes, denoting the interest in maximizing local cache efficiency in potentially highly mobile network environments.

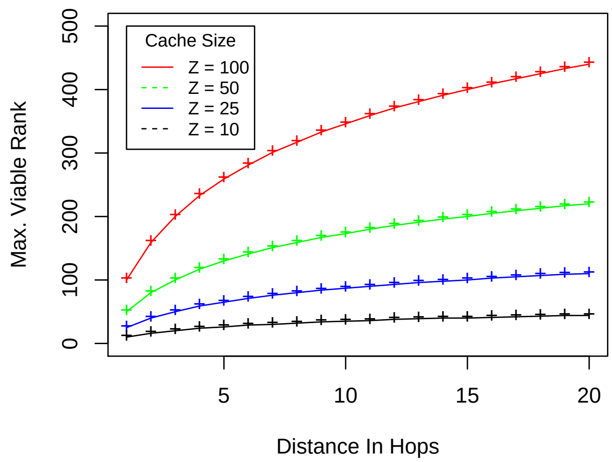

Figure 7.2: Replaceable Rank Threshold Evaluation for Z Cache Sizes

Moreover, the gain from increases in distance display a similar result to those observed in the work presented in Chapter 5. In both cases, increases cache sizes, and subsequently increased variability, had limited capacity to improve cache hit rate. This is the result of the low probability of request for items with high ranks and is the motivation for limiting cache variety. As an example, consider the item with rank #250, which has a likelihood of being requested of approximately 0.049%, making it a less preferable choice for cache inclusion. Because content items with high ranks are unlikely to be requested, caching such items is not necessarily beneficial, whereas caching items with lower ranks has potential to reduce network load significantly via responses from closer nodes. This way, increases in cache size have limited benefit, further reduced by the need for caches to operate line speed, making maintaining large caches not viable. Hence, given the diminishing returns and other limitations of cache size reliance, for the remainder of the evaluations, we consider $Z = 50$.

Increased cache diversity can be achieved by increasing the λ distance impact factor without the need for increased caches. Evaluation of the influence of λ is assisted by the results displayed in Fig. 7.3. The graphs display a similar behavior to that of increased caches, albeit in a way that is programmable. While for many scenarios, increased cache might not lead to higher hit rates, allowing for selective filtering can yield positive results in specific circumstances. For the given graphs, the cumulative probability solely for the

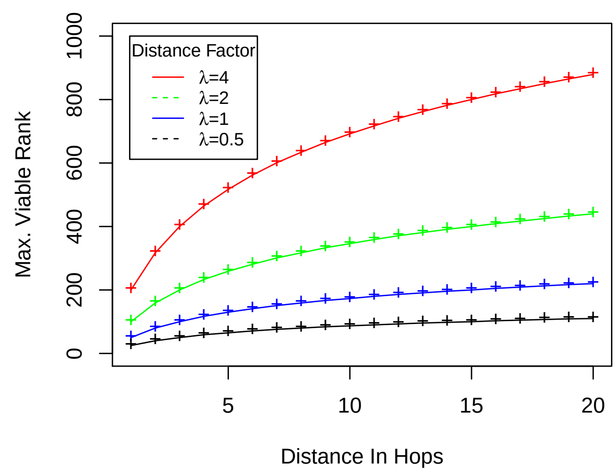

Figure 7.3: Replaceable Rank Threshold Evaluation for λ values

lowest ranked items is 55% while for the $\lambda = 1$ the potential top 250 candidate items have a cumulative probability 74.5%. For $\lambda = 4$, the 850 potential candidates have a cumulative probability of 89.54%, again denoting the diminishing return in increased variability. Conversely, λ can be set to restrict cache variety, making networks more selective. Increase selectiveness is especially interesting for scenarios with high popularity skew or small caches. By restricting cache variability, the policy can be made to operate on a strictly popularity-centric format. The evaluations made so far are based using the precise content popularity values, corresponding to a state where nodes operate with global knowledge. However, partial knowledge has a significant effect on the efficiency of cache solutions, leading nodes to operate on limited and potentially incorrect information. As cache sizes are usually disproportionately smaller than the available content library, item perception of the content space is relatively limited, meaning that items with infrequent request patterns are hard to tell one from another. This is yet another motivator behind the inclusion of the λ factor. One of the main consequences of knowledge limitation is a limited sample size for object requests. Items with a smaller request frequency (higher ranks) have increased level of uncertainty attached to them. As higher rank items are less likely to be requested, nodes face difficulty when discerning item request frequency for these infrequently requested objects. Hence, it can be beneficial to avoid cache operations when both cached and incoming items request frequencies are low, as the replacement is likely needless. At the same time, to allow items with few requests but inherent higher probabilities from being neglected due to relative frequency obfuscation, storing previous frequency data is fundamental. This is necessary to avoid having caches not reflecting the network state as well as to increase calculation precision. If only the servers were to log requests, cache usage updates would not propagate and nodes would operate with significantly less knowledge. In this case, the network would become dependent on server knowledge of request frequency, limiting intermediate cache efficiency. As partial knowledge affects the distance factor parameter, its value can be manipulated depending on specific network properties to compensate for mobility and network density. These consideration, however, are outside the scope of the work presented in this Chapter, and is listed under the list of future research efforts presented in Section 9.2. Regardless, based on the numerical evaluations, we are confident in the efficacy of the UG-Cache and have proceeded to perform additional simulation-based evaluation of the technique, presented and discussed in the following section.

7.4 Evaluation

This Chapter's main contribution is an implicit coordination policy for cache insertion for ICN VANETs. In this section, we present a comprehensive set of simulations designed to evaluate the performance of the proposed algorithm. With this experimental evaluation, our intent is to further illustrate the merits of the proposed cache insertion coordination policy. The simulation set was carried out using the OMNet++, SUMO and Veins simulation tools [100, 102, 101]. Simulation environment setup is based the model discussed in 5, consisting of a downtown area road network with dimensions $2km \times 1km$. Simulations are run for 600 seconds preceded by 100 seconds of setup time, during which vehicles are instantiated and dispersed throughout the network. Vehicle inclusion and movement patterns are based on road density and road-type preference and were configure to replicate real-world characteristics. Vehicle density is fixed during simulation, with vehicles that completed their routes being re-routed upon reaching their target. The setup configured allowed for the simulation to maintain a "ready state", maintaining cache information throughout the simulation.

The simulation scenario considered contains a total of 150000 thousand individual content objects, grouped into three categories with 50000 items each. We chose to consider multiple simultaneous categories to emulate scenarios with co-existing application spaces. The popularity distribution of items within each category is characterized by the

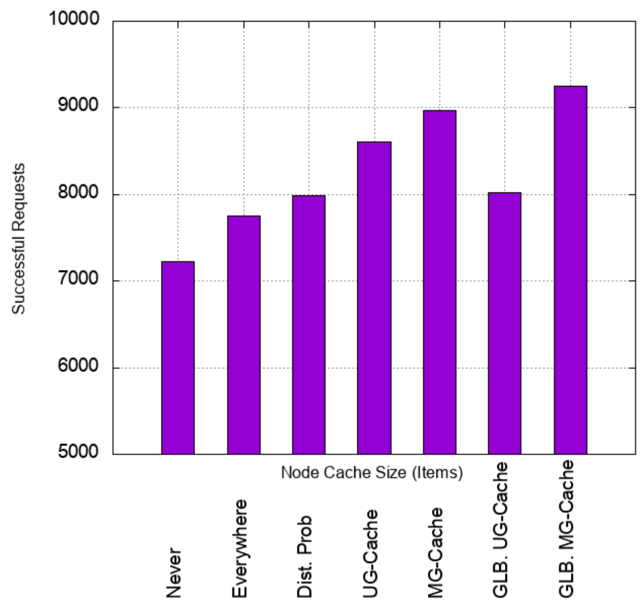

Figure 7.4: Total Requests Fulfilled

Zipf-like model described with $\rho = 0.8$. Each content object was set to have a size of $100kbytes$. Based on the numerical evaluation performed previously, caches were configured to store a total of 50 objects at any given moment. Networks were configured to have a constant volume of 300 vehicles, 10% of which were configured to be content producers, serving as servers that replied to all content requests, while all other nodes operated clients and as intermediate nodes. Request generation followed a Poisson distribution with average inter-request timers set to $0.25s$. The proposed UG-Cache solution is evaluated against the works listed previously, namely: Cache Everything Everywhere, Never Cache, Distance-Probabilistic. We also compare UG-Cache to a variation of itself where global popularity knowledge is available. This evaluation is made for both average UG-Cache and MG-Cache. This inclusion is made to denote the influence of partial knowledge in content popularity perception. Following is the discussion of the results obtained.

The main goal of UG-Cache is to increase cache hit ratio, reducing hop distance and, consequently, network load. We evaluate the proposed UG-Cache and MG-Cache solutions by comparing them to existing works in terms of Content Delivery, Cache Replacement, Latency and average hop count. Simulation results are displayed for the averages of over 30 iterations per configuration and where 95% confidence intervals dis-

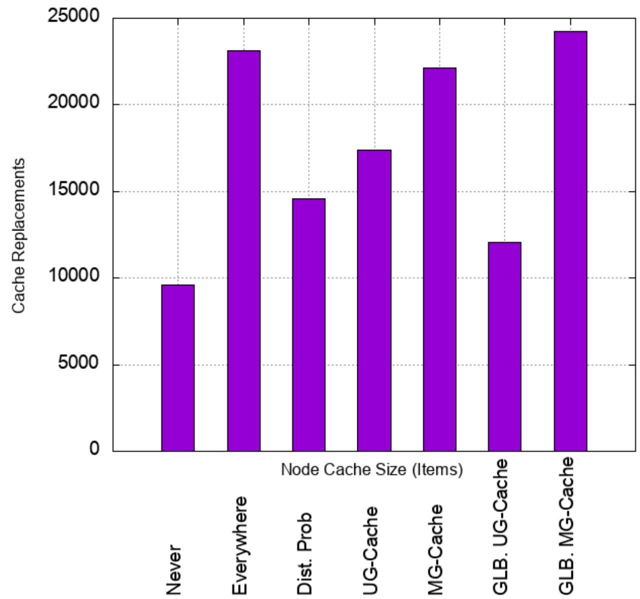

Figure 7.5: Total Cache Replacements

play negligible deviation.

In Fig. 7.4 we present total successful requests for all approaches evaluated for $\lambda = 1$. From the results, it can be observed that both UG-Cache variations lead to increased content delivery. When compared to the alternative schemes, both UG-Cache and MG-Cache yield positive results. Application of MG-Cache leads to an increase of up to 25% in the delivery rate while UG-Cache leads to increases of approximately 17%. Comparatively, relative to the Dist. Prob. policy, gains are of 15% and 8% respectively. Global knowledge variations of the proposed solutions, on the other hand, have limited increase when compared to existing solutions. The Global UG-Cache variant does not lead to any improvement, while the Global MG-Cache policy results have a 5% gain over existing solutions. These results are attributed variation in how knowledge affects the solution. For UG-Cache and MG-Cache, the partial knowledge nodes have when sharing information with their peers lead to more frequent cache replacements, as seen in Fig. 7.5. The increased replacement rate in these solutions results in increased content pools, which lead to improved content delivery. However, unlike alternate approaches, cache inclusion is dependent on utility potential, leading to overall increased cache hit ratios. In fact, MG-Cache performs almost as many replacements as the CEE policy, yet increases total deliveries by as much as 16%. From a social standpoint, cache diversity is more balanced, ensuring that popular items are available in caches while increasing variety. In contrast, the Global UG-Cache policy variant leads to a significantly less varied cache. This can be observed in terms of the cache replacement count for the policy, which is nearly as low as the results displayed for the policy with no intermediate caching. Yet, comparatively, the Global UG-Cache does allow for an increase in request, reinforcing the benefits of popularity awareness. The limited cache diversity, however, restricts content delivery whenever content servers are not reachable, leading to lower delivery rates compared to UG-Cache. Based on this we can conclude that cache variability dependent on popularity yields positive results, with partial knowledge having limited impact and allowing for increased variability, with potential additional gains.

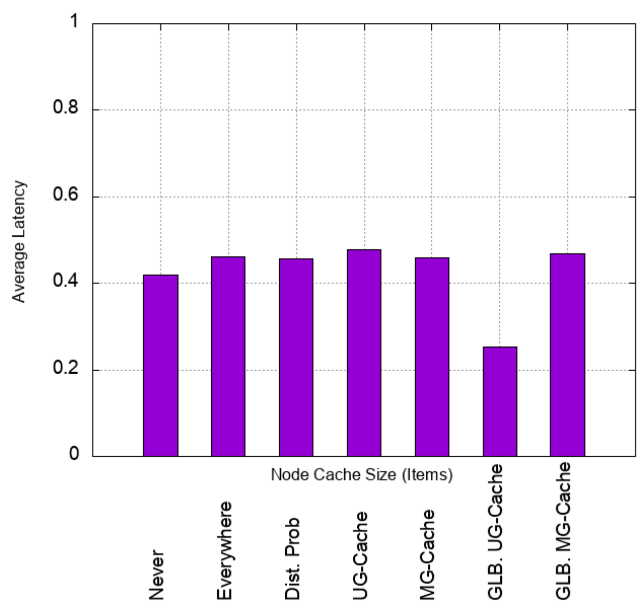

Figure 7.6: Average Latency Results

The gains promoted by UG-Cache and MG-Cache come at negligible latency costs. Fig. 7.6 displays the average request response latency. As can be observed from the graph, nearly all policies have similar latencies regarding content delivery except for the Global UG-Cache variant. This is the result of the low cache insertion rate targeted towards high popularity content. Global knowledge significantly reduces cache variety, increasing duplication and leading the fewer items that are cached to be served faster. While both Global variants have been included as reference, such behavior is a clear example of the issues that can potentially arise from use of popularity and a motivator for the inclusion of variant schemes such as MG-Cache and the λ impact factor. Variation in latency between UG-Cache and MG-Cache to alternative solutions are around 7%, result of a slight increase in distance to content. Increased latency is more notable in the Global MG-Cache variant results, with a 15% increase in latency result of the further increase in cache variety. Overall, results are satisfactory and pinpoint the clear benefits of the proposed policy, which is poised to have additional gains over policies given it was compared to solutions operating under random replacement rules, as is commonly done in the literature. Regardless, even under optimal conditions, alternate solutions are less capable of promoting efficient content delivery and network resource use.

7.5 Closing Remarks

In this Chapter, we proposed and evaluated a cache insertion policy for ICN VANETs named UG-Cache and its variant, named MG-Cache. Both solutions were based on the concept of estimation of item cacheability through use of a utility function that considered both popularity and cache distance. Estimation of item usefulness was evaluated numerically in terms of absolute probabilities to determine potential item eligibility, denoting a need for balance between cache variability and maintaining highly requested items to increase delivery rates. Simulation results confirmed the benefits of the proposed approach with benefits from UG-Cache and MG-Cache being complementary, allowing for a potential for application of both solutions as well as on a potential integration of both alternatives.

Chapter 8

Social Pre-caching of Location-dependent Requests

The contents of this Chapter are divided as follows. The Chapter begins with a contextualization of location dependency in caching, presented in Section 8.1. In Section 8.2 we present the context and overview the literature in location correlation in networking. Following, in Section 8.3 we further contextualize Provider-based Location Correlation, including a preliminary simulation demonstration of the predicted limitations. From this discussion we present our proposed solution in Section 8.4. Then, in Section 8.5 we evaluate the proposed solution, followed by our final remarks, made in Section 8.6.

8.1 Introduction

In-network knowledge production and ad-hoc dissemination are among the many defining properties of Vehicular Networks. Being part of Intelligent Transportation Systems (ITS), VANETs have access to multiple content production sources other than the Internet. Increases of in-vehicle sensor deployment further expand information generation capacities, allowing vehicles to serve their peers with a greater variety of applications. A natural consequence of knowledge exchange is the emergence of service exchange. Service exchange can be explored for multiple applications, including those outside the initial scope of vehicular communication. Initially studied in terms of emergency and safety related services, vehicular networks have potentials in transit and infotainment. Ultimately, the mechanisms for information dissemination and exchange will continue in relevance in the context of vehicular networks.

Caching is one such area that can provide significant improvements to the overall service quality. By adapting to patterns unique to a region, nodes can improve cache insertion/eviction potentially increasing cache hit rates and, consequently, reducing the average latency and network load. Further improvements can be made by collaborative social cache coordination. Information-Centric Networking (ICN) is an alternate approach to communication from host-to-host that is tailor-made for application-centric systems and allows for universal naming and increased cache applicability and coordination. Both explicit and implicit coordination policies can be taken advantage of in ICN to increase cache variety while maintaining highly popular content objects readily available. Notably, however, there are few works on location-dependent content caching and social caching in mobile networks, with even fewer works focusing on vehicular networks. This is partially explained by the ever-evolving nature of wireless networking and recency of vehicular networks and information-centric networking. Given the tremendously ample scope of VANETs, we see an opportunity to address an issue faced by many applications with a focus on traffic information dissemination within the context of ICN.

8.2 Literature on Location-dependent Caching

Locality is an important property of networking with potential for increased research, specifically in cache design. This argument is made considering vehicular networks, where due to mobility and potential applications, locality is increasingly important. While location-dependent caching has been explored in wireless networking, it has been over a decade since the last set of works focusing on the topic [114, 69]. Hence, given the inherent difference in nature between ICN and host-to-host communication, additional research is required. Moreover, given the unique properties of vehicular networks, ICN-VANET caching is a topic with potential for research. In general terms, location dependent caching is a subset of cache policies and schemes designed to optimize caching based on the spatio-temporal properties and patterns of either the nodes or the content itself. Node mobility, for example, can be taken advantage in implementing cache placement and routing strategies [113, 33, 64]. Caching policies based on content spatio-temporal properties, however, has had fewer works dedicated to it. Wang et al. denote the difficulty of determining and exploiting patterns [115]. While temporal correlations have been investigated by works using frequency estimation, little work has been made in location-correlation. Location correlation, or locality, is generally defined as an asymmetry in spatial distribution for content requests with location-dependent variation in the probability of content being requested. However, we consider this definition generic and allowing for two conflicting interpretation of locality, one in which the requester location is the influencing factor and one in which the provider location is the governing property. In vehicular networks, both correlations between requester locations and between provider location are contextually plausible, as will be discussed next.

Requester-based Location Correlation (RLC) is the approach most commonly considered by works [62, 116, 115, 69]. This interpretation establishes that the position of content requesters (client nodes) affect the distribution of frequency of the requested objects (content popularity curve). Simply put, the probability that an object $\mathcal{O}$ is requested varies depending on the position of vehicles. Hence, the global popularity distribution of a content object by itself cannot fully describe the likelihood of the said object being requested. The spatial distribution of popularity for any given object may be skewed with locations of increased or decreased popularity. Contextually, the skews in popularity are explained by interest locality, commonly determined by social interests. This behavior can be observed in social media, with spatio-temporal proximity affecting content request patterns [117]. In their work, Kassio et al. analyze a large set of social-media posts in metropolitan areas and conclude that certain topics are inertly associated with specific locations, displaying fluctuations in relative frequency in their associated regions. Said frequency changes were consequences of either the properties of the locations such as their purpose (e.g.: events at sports complexes) or due to correlated characteristics of the users in those locations (e.g.: routines). Previously, researchers have considered traveling and traffic applications as being susceptible to locality [69]. In their work, Xu et al. consider that location-dependent requests are bound to the physical location of their users and the effects of node mobility. The examples presented by the authors relate to travel information and traffic information requests, where nodes inquire about information in the region they are in. In this work, content objects are associated with specific cells and client mobility can invalidate locally available content depending on its validity scope. This is due to content ceasing to be "nearby" as nodes travel away from a specific region. In this case, nodes explicitly do not request objects outside their region. Another work studies cooperative caching in face of localized request patterns [116]. In their work, Sluijs et al. study the effects of locality in a ring-shape wired network. The authors denote that given correlated user interest, cooperative caching solutions can outperform proactive replication. The interest model proposed by this work specifies item popularity in terms of a Zipf-like distribution multiplied by a $P_{Normal}(y)$ personal content item interest probability function where y is the index of the requesting node. According to the work, higher values of σ make the function more uniform, with neighboring nodes having a higher probability of requesting a given item. The approach taken by this work reshapes the popularity distribution function, reducing the popularity of its most popular items while maintaining a log-normal tail. From these examples we observe that all works focusing on location dependency establish some sort of *requester-centric* correlation but with particularities in their modeling.

Provider-based location correlation (PLC) an alternative interpretation to location-dependency that has had little research focus. Unlike Requester-based correlation, in PLC, the location of content providers is correlated rather than that of its clients. This difference has significant consequences for content delivery. In RLC, as clients are close to one another, clients may obtain objects from nodes next to them who have requested the same objects previously. In PLC, however, nodes are less likely to benefit from content cached by previous requesters. Moreover, as content providers are near, all requests end up being directed to the same area, leading to increased average distance and increased load in the location where the content providers are. Relating to the previous example of traffic information, considering the navigation context, vehicles generally are interested in knowledge on the direction to which they are traveling to. In this example, the content servers are the vehicles close to the areas of interest. In locations such as urban “artery” roadways and larger intersections, may face significantly increased requests compared to the rest of the network. The need to tackle the issue of provider-based location correlation from an intermediate caching perspective becomes clearer considering the potential improvements, including improved network load distribution, reduction in the average distance to content and improvement to overall communication quality, all result of increased efficiency cache management and coordination for location correlated content. Following, we further discuss Provider-based correlation and make the case for the need for improvement in content delivery mechanisms targeted towards this kind of application.

8.3 Location Correlation

To address the limitations imposed by location correlation, we must first discuss how this property manifests itself and what are the possible consequences. Following, we provide additional discussion on Client-based Location Correlation to introduce how we modeled this solution and then perform baseline simulation-based evaluation designed to denote the limitations of provider-based correlation when compared to client-based correlation.

8.3.1 Provider-based Location Correlation

Provider-based Location Correlation (PLC) is an interpretation of location dependency in which the provider location is the informing correlation factor and one that has yet received its deserved research focus. The relevance of traffic information dissemination in vehicular networks motivates its selection as the context for the in-depth discussion of PLC. PLC may also be observed in any other context where content provisioning is tied to real-world data collection such as weather, live-feeds and even services such as parking availability. In the case of traffic information requests, knowledge of traffic condition is initially only available to those who are in the specific region. As content provisioning is tied to a geographical location, request dissemination and content provisioning are hindered by increased distances. Techniques such as Geographic Routing can be used to mitigate this issue, restricting broadcast storms and reducing network load. However, without introducing caching policies targeting PLC content objects, distance to content is likely to remain high. We argue that selective preemptive dissemination and intermediate caching of popular location-dependent content is fundamental to increase the efficiency of content delivery while simultaneously distributing network load away from the original content source location.

To understand the challenges in content delivery relating to PLC content, we selected to perform a set of simulations, designed to denote the limitations of PLC when compared to non-geographically correlated content as well as to RLC requests. To expedite the simulation process, we take advantage of a previously validated simulation environment that emulates real-world conditions described in Chapter 6. Basic network and mobility configurations described previously were used as a basis to generate a new set of simulation scenarios and results within the requirements of the scope of the work in this Chapter. All other properties, and protocols were either re-implemented or are novel to the work in this Chapter, being described next. Regarding the specific properties of the simulation scenario, its boundaries delimit a $2km \times 1km$ cross-section Ottawa's Downtown Area. Vehicle volume is set to 300 vehicles, with a density of 25 vehicles per kilometer of linear paved road. Simulation is preceded by an initial 100 second startup time necessary for vehicle instantiation and dispersal, and is followed by 300 seconds of communication simulation. A total 5000 content objects are available for request, with item popularity varying according to a Zipf-like distribution with $\rho = 0.9$. All content objects occupy a predetermined $100Kb$ of caches, and the caches themselves are sized at $5Mb$, for a storage of potential 50 objects. Locally, nodes request objects on average every 0.5 seconds, and non-cached requests being forwarded to the on-board network interface. To establish location patterns, we opted to design novel location correlation models as given the limited scope of in simulation scenarios modeling in related literature, as discussed previously. Regarding PLC contents, we divide the geographical simulation region into square grid sectors of size κ meters. Sectors are then randomly assigned ranks, which are mapped to their Zipf-like request probabilities. Hence, the same mechanism used for non-correlated requests is extrapolated to PLC request distribution, ensuring similar probability distribution. For this work, we consider $\kappa = 20$, again resulting in a total 5000 sectors which represent an equivalent number of content objects. We also consider the maximum distance up to which vehicles can reliably provide information regarding a sector, defined as β, and is set to $\beta = 200$ during evaluation. Following we discuss RLC and its particularities as well as how this mainstream interpretation is considered for our evaluation.

8.3.2 Requester-based Location Correlation

The simulation setup acquired for this evaluation is flexible but did not contain explicit models for PLC and RLC requests. As discussed previously, there is are logically consistent mechanisms in the related literature available to study RLC. However, as mentioned, the works on location-correlation, generally define restrictions and other particularities that limit their applicability as general vehicular ad-hoc networking solutions. Given, that to the best of our knowledge, current literature lacks the definition of a model that allows for the evaluation of RLC in a realistic scenario, we propose that related work from social media big data analysis be used to enable the necessary study. We take advantage of the data obtained by Kassio et al. [117], kindly made available for this portion of the research. The dataset provided contains over 1.5 million geolocated social media messages with associated hash-tags from the New York City region and are temporally distributed over a period of 1 year.

To improve compatibility between the dataset obtained and existing simulation configurations, the approximate 90 thousand unique social objects were batched into subset group via a 1/18 floor function. The resulting dataset has approximately 5040 items and maintains the $\rho = 0.9$ Zipf-like distribution characterization observed in the original dataset, as observable in Figure 8.1. Similarly, the temporal distribution is compressed to the simulation time. From the batched dataset we projected the central Manhattan island area into the simulation area, with the resulting spatial message distribution

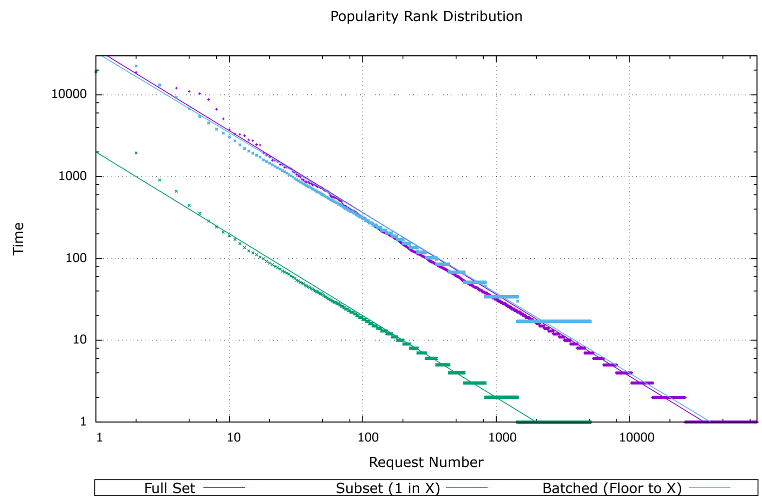

Figure 8.1: Frequency Distribution of Content Objects in Social Media Data

presented in Figure 8.2. The projection-mapped data contains approximately 500 thousand unique messages. Note that while message distribution is not perfectly symmetrical among the mapped region, that is not an issue, as spatial distribution is not homogeneously distributed in original data. Moreover, messages need to be assigned to vehicles during simulation, further reducing the need for absolute symmetry. Request assignment is performed by a central simulation controller, which selects vehicles close to the original request position. Vehicle selection prioritizes proximity but considers frequency distribution amongst potential clients during selection. The maximum viable assignment distance, also informed by β is set to 200 meters, but in this context refers to the distance from the original request position. The spatial requests assignment map and resulting communication patterns observed for RLC are comparable to the scenarios without spatial correlation. With a RLC model in hand, we can proceed to the proposed preliminary evaluation.

8.3.3 Preliminary Evaluation

In this preliminary analysis, our focus is determining the impact of Provider-based location correlation (PLC) in network efficiency. Towards this, we present the following simulation results, executed on all three distinct location correlation models discussed. Figure 8.3 presents the average delivery rates for a vanilla NDN-based network over a variety of cache policies for both PLC, RLC and a scenario with no spatial correlation.

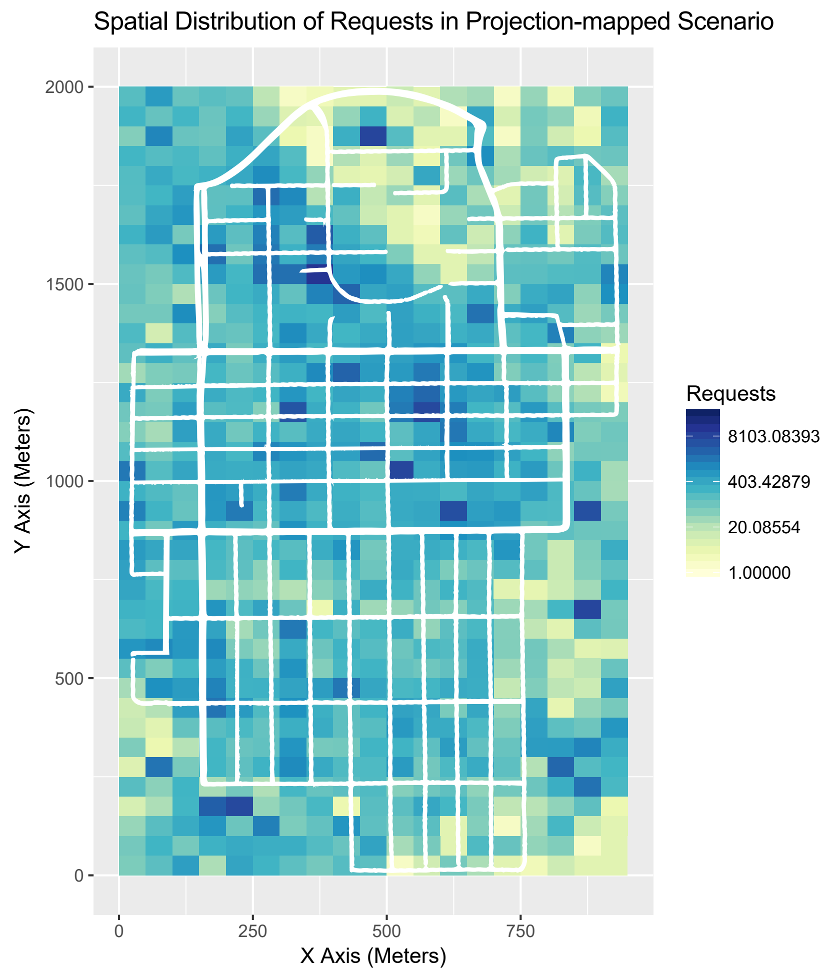

Figure 8.2: Projection-Mapped Spatial Distribution of Social Media Requests

The considered intermediate cache insertion policies are namely, Never Cache (No Intermediate Copies), Leave Copies Everywhere (LCE), Leave Copy Down (LCD) [71, 68], and UG-Cache as well as its variant MG-Cache, both introduced in Chapter 7. As can be observed, for every cache policy, both randomly distributed requests and RLC configurations yielded positive results with high delivery ratios. In scenarios with Provider Location Correlation, delivery rates are reduced as much as 50% with further decreases in the scenario where intermediate caching is not present. While most cache policies yielded improvements over the policy in which no intermediate caching, results remain disproportionately lower than those where there is no provider correlation. These results strongly reinforce the previous discussion regarding the need for improved content

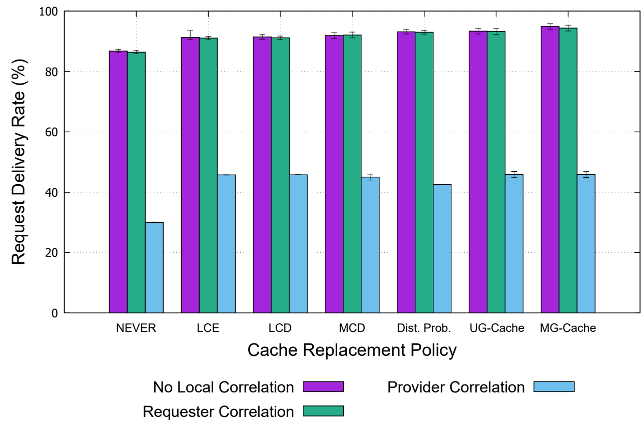

Figure 8.3: Delivery rates for a "Vanilla" network without PLC specific Policies

distribution throughout the network. Notably, additional simulations with increased vehicular density lead to further decreases in location-based content delivery. This was observed to be the result of increasing competition for network resources and diversity of caches, which culminate in a reduction of up to 75% in delivery rates for scenarios with twice as many vehicles. These results reinforce the previous conclusion, which denotes the challenges in content provisioning associated with asymmetrical content placement previously discussed in Chapter 7. Most cache policies displayed similar efficiency in contrast to the configuration with no intermediate caching. This is relevant as it shows that even when not designed explicitly for the challenges of PLC, intermediate caching is beneficial, a fact which we take as a starting point for the proposed solution, presented next.

8.4 Proposed Solution

The main issue tackled by this work is the delivery of provider-centric location-dependent content. As denoted by our preliminary evaluation, asymmetrical availability content objects to restricted areas impose challenges in terms of content dissemination and can significantly reduce communication quality, result in an increased distance from clients to content providers. To improve content delivery, we propose an approach based on preemptive content caching. The mechanism proposed is based on controlled distribution and caching of items in a reserved portion of each nodes' cache. We take advantage of the prefix scheme proposed in our previous work, where the naming scheme is structured based on application class and context, allowing for explicit classification of services as geo-located, as discussed in Chapter 6. With this knowledge, whenever a node receives a request for an item with location-dependent patterns, it stores this log in a separate table. This applies to both vehicles that can fulfill the request via local content producers and sensors as well as by vehicles that have previously cached the requested content object. By maintaining records of overheard requests in their given region, content providers gather knowledge of the frequency distribution. The accrued knowledge is shared by content providers in a cyclical process. The information shared is used by neighboring nodes, informing their pre-fetching decision regarding the objects publicly disclosed. PLC items cached by this mechanism are stored in a reserved portion of the cache and can only be replaced by other elements obtained similarly obtained by explicit pre-fetching requests.

For every incoming *interest* message, if the content object requested is *available* locally and belongs to an application with spatial correlation, interest statistics for the requested item are recorded. Statistics are maintained in a weighted exponential sliding window, with the frequency of each item requested being logged separately for each temporal slot. The forgetting factor weight of the sliding window is defined as α. Each period contains the statistics logged between two beacons, and n periods are kept in total.

Based on the statistics collected, prior to beacon transmission, nodes perform popularity evaluation on the temporal interests recorded. Content prefix and relative frequency for the most popular object are broadcast as meta-data along with beacon messages. The procedure, designed with the purpose of selecting popular items while limiting beaconing of unnecessary data works as follows. For each item cataloged, the weighted average of request occurrences is calculated using the forgetting factor α. The item with the greatest frequency is then compared to the minimum viability threshold λ, necessary to limit the potential overhead generated by pre-fetches.. If the selected item has frequency above λ, its statistics are broadcast to the network. Whenever a node also located within the geographical location a PLC request overhears a content cache suggestion, they perform two actions. First, upon receiving the advertisement, they log the popularity statistics locally in a vector that lists all overheard information during an ongoing cycle. Then, prior to their own beacon advertisement, the information vector cataloged is computed into the final list used for beacon object selection. The last step of the process for all advertising nodes is the maintenance of the temporal series, shifting the sliding window for the next data collection cycle.

For nodes outside the PLC region, whenever a node receives a beacon that contains

Algorithm 1: Cache Decision Algorithm

```
Result: Cache Insertion decision for object 𝒪
Min ← INT_MAX
Average ← 0
Near ← NULL
ϕ ← 𝒪.frequency
foreach Object I in PLC_ Cache do
    θ ← ℐ.frequency
    Average ← Average + θ
    if P_f(θ) < P_f(Min) & DistanceTo(I) < δ then
        Min ← θ
    end
    if Distance(I,𝒪) < β & I < Near then
        Near ← θ
    end
end
Average ← Average / Size(PLC_ Cache)
if Near < ϕ & Near < Average then
    PreEvict(Near)
else
    if P_f(ϕ) ≥ P_f(Min) then
        PreEvict(Min)
    else
        Exit()
    end
end
Request(𝒪)
```

an advertisement of seemingly popular PLC object it does not currently have in either cache spaces, said node evaluates the potential for the item compared to its current cache status. The decision process takes into consideration the current state of the cache section allocated for PLC content as well as its regularly cached items. If there is space available in the PLC cache, the vehicle proceeds with a request for the advertised content object. Otherwise, the incoming cache suggestion is evaluated using Algorithm 1. The algorithm evaluates the reported popularity of the incoming suggestion the statistics provided by previous advertisements as well as local statistics of requests for said items. The decision considers the following equation, designed to estimate the potential positive impact of inserting object $\mathcal{O}$ into the PLC sub-cache:

$$P_f(\mathcal{O}) = P_f(\phi, d, t) = \frac{\phi}{2^{t-(d+1)}} \tag{8.1}$$

where ϕ is the advertised frequency for the object $\mathcal{O}$, t is the number of beacon cycles since object $\mathcal{O}$ was initially observed/cached and d is the distance from the advertiser. Equation 8.1 is evaluated against items in the PLC cache. The item with the lowest recorded value for I_f is selected and marked for eviction and a special PLC interest message is broadcast towards the original beacon broadcaster. An advertised item might also be selected for insertion by clients should it be more popular than another object mapped to the same region, where the maximum regional range, as mentioned previously, is β, given in meters. In this case, if the dedicated PLC cache is full, the less popular object is pre-selected for removal, to be executed upon the node receiving the requested PLC object. This is necessary to avoid local redundancy, and functions given that popularity informs general cache placement probability, with PLC observations having more general likelihood of requests when compared to locally relevant objects, which are potentially served by neighbor nodes. Local redundancy amongst caches is avoided to reduce local redundancy of objects. This is because popularity already informs cache placement, meaning that should multiple items be locally relevant, request frequency can inform overall cache decision for other objects, with the PLC Object that has the greatest likelihood being preferable over the least relevant. Upon overhearing advertisements for GPS objects, should a client decide not to cache the object, they also act and log the overhead statistics. Should the node have already observed a suggestion for said PLC object, it updates its stored statistics to contain the highest request count and shortest distance to object. Then, similarly to the PLC content providers, the listening node decides if it should re-broadcast the updated information along with its beacon. Note that nodes consider the δ maximum broadcast hop distance when deciding whether to broadcast a message. Algorithm 1 presents the decision process in greater detail, elements being expressed in terms of the objects being evaluated, omitting the individual properties to simplify notation.

Upon receiving the PLC Interest request, the server proceeds with content provisioning. To limit replication, upon receiving a request, the provider halves the recorded statistics for that specific item. Any subsequent pre-cache request will be accompanied by the minimum threshold established at the client side and with the updated information the content provider can make an informed decision on whether the request still warrants a response, or should it be ignored. In a negative case, nodes do not proceed with another request and ignore the unresolved request, as they can potentially re-request the object upon receiving another advertisement. This is necessary both to avoid said excessive redundancy in PLC caching but also to limit the overhead caused by the pre-fetching mechanism. Only a delivery attempt in the reverse path has recovery mechanisms, as that is the expected behavior for content delivery. Lastly, note that both interest messages logging and subsequent preemptive cache requests are performed off-line and do not delay interest responses. Therefore the proposed solution allows for the more complex cache analysis operations when compared to in-line request response. From this description, we proceed to the evaluation of the proposed model.

8.5 Performance Evaluation

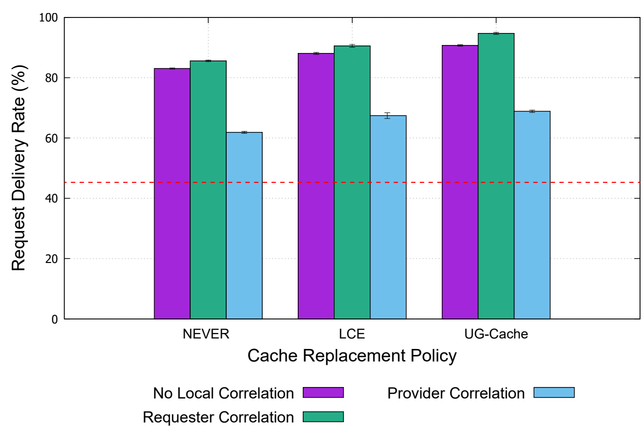

Figure 8.4: Delivery Rates for a naive PLC Policy

To Evaluate the proposed pre-caching policy, we consider the same scenarios and configurations described in Section 8.3.3 both for the sake of consistency and to reduce the need for additional descriptions. Regarding the policy, as proposed in Section 8.4, we consider for this evaluation the following parameter values: $\alpha = 0.75$, $\beta = 200$, $\lambda = 50$, $\delta = 10$ and $n = 10$. We also set the value of the reserved PLC sub-cache to $\frac{1}{10}$ of the total cache space, equating to 5 objects. From this setup, we performed a series of simulations.

Figure 8.4 presents delivery rate statistics for three most relevant cache insertion policies grouped based on the content correlation model. Compared to preliminary results presented in Figure 8.3, we observe an improvement in total content delivery of between 20% and 50% for PLC content when compared to the scenario lacking any policy. The resulting difference varies depending on the associated cache insertion policy. Specifically, for the scenario with no intermediate caching, the increase is more significant as it introduces an intermediate cache coordination policy where there was none. Naturally, this speaks to the benefits of intermediate caching, and of ICN by extension. More importantly, it denotes a significant relative reduction in the difference in delivery rates from a policy with no intermediate caching and those with one from 30% to 3%. Moreover, by itself, it speaks to the efficiency of the proposed solution, which without intermediate caching, can significantly improve delivery. Notably, delivery rates remain lower than those of non-correlated or RLC scenarios even if significantly improved. This further speaks for the necessity of research focus on Provider Location Correlation.

Additional results from the same simulation set denoted an increase in average hop distance, with an increase from an average of 1.2 to 1.63, including requests fulfilled by local PLC Caches. This is the result of an increase in cache diversity promoted by the policy proposed. While the average distance to content increased, this was necessary to enable the increased content delivery. Moreover, overall network load, which displayed an increase of up to 21.9% in PLC scenarios with no policy, has marginal differences in load compared to the non-PLC scenarios whenever the proposed policy is deployed. From our observations, further increases in content delivery are likely only achievable via increased coordination or efficiency in item-selection for pre-caching.

8.6 Concluding Remarks

In this Chapter, we have brought light to a limitation faced by many applications in Vehicular Networks. We introduced a simulation model to evaluate both modes of location correlation and used said model as part of a preliminary simulation. The simulation evaluation further supports the previous observations, displaying a significant gap between content delivery in applications with Provider-based Location Correlation (PLC) and those with either Requester-based correlation or those where there is no correlation. From the observations, we proposed a policy targeted towards PLC content which leads to improvements in content delivery without hindering other quality metrics. The results obtained denote benefits of solutions targeting PLC and the need for further studies, that, hopefully, can lead to PLC content having delivery ratios on-par with its alternatives.

Chapter 9

Conclusions and Future Work

This chapter is divided as follows: Section 9.1 presents the summary of this thesis. Following, in Section 9.2 presents a series of research opportunities relating to the topic covered by this thesis.

9.1 Thesis Summary

This thesis focused on Information-Centric Networking in the context of Vehicular Networks and its overall contributions related to the introduction of a information-centric architecture and caching algorithms designed to enable and improve service exchange in ICN-VANETs. The methodology employed focused on deriving research opportunities from literature review and subsequent proposal, design and evaluation of novel elements. Chapter 1 presented a brief introduction to the topic and provided the context necessary to present a research statement focused on enabling service exchange in ICN-VANETs. Following, we presented the current research trends and state of the art in ICN-VANETs Chapter 3. As part of the discussion on the existing literature, we denoted challenges relating to content naming, routing, content discovery and caching. From the literature review, we observed a significant limitation when it came to the discussion of public transit as part of vehicular networks, which we felt required further investigation. In Chapter 4 we presented an analysis and discussion of the roles of public transit as part of vehicular networks, potential applications and challenges, as well as a brief evaluation that justifies the discussion made. This evaluation was made in the form of trace-based analysis, which denoted significant potential for city-wide road network coverage.

From the comprehensive contextualization performed and the research opportunities observed, we proceeded to the introduction of a communication architecture designed specifically for ICN-VANETs. In Chapter 5 we introduced *SEVeN*, an Information-centric architecture designed with the intent of facilitate the implementation of services within Vehicular ICNs. This included the definition of specialized service managers and orchestrators tasked with implementing the functionalities necessary for service exchange. SEVeN enables efficient delivery of content model while ensuring that priority services can maintain functionality in face of high network load as well as achieving the goal of promoting content sharing via coordinated access to the network resources. A novel naming-structure and content prioritization model designed to enable service classification and QoS management were also introduced as part of the definition of SEVeN. The proposed architecture served as a boiler-plate that allowed us to overcome challenges impeding the implementation of service exchange in vehicular networks and serves as an example to foment attention to the topic of collaborative vehicular networks.

Following, we focused our efforts on discussing caching, a fundamental component of computer communication with significant potential for improvement in ICN environments. To better understand caching in the context of ICN, in Chapter 6, additional literature review on caching was performed followed by a preliminary performance evaluation of applicable solutions. We demonstrated that in the context of VANETs, mobility is paramount and that the spatial distribution of content is a key factor in caching efficiency, pinpointing techniques that can be incorporated into future solutions. From this discussion, in Chapter 7, we presented a pair of novel implicit-coordination policy variants designed taking into consideration the unique challenges of vehicular networks. The core concept that informs both UG-Cache and MG-Cache policies is the estimation of content utility, defined in terms of content popularity and cache distance. By taking the likelihood of access into consideration, the proposed policies avoid excessive duplication, which as observed previously, is detrimental to system performance. Simulation performance evaluation confirmed efficiency of the proposed cache policies and the need to balance cache variability while increasing spatial content distribution. Finally, in Chapter 8, we brought light to interest locality, a limitation faced by many applications in Vehicular Network. To build upon our previous research and enable the study of this issue within the context of VANETs, we introduced a simulation model to evaluate location correlation. We described both potential interpretations of location correlation and performed a set of simulations that demonstrated the challenges faced in contexts where provider location is the determining factor in request fulfillment. We then introduced a pre-emptive cache coordination policy targeted towards this type of content, which lead to significant improvements in content delivery without hindering other service quality metrics.

9.2 Future Research Opportunities

Future research work for consists of further investigations in service exchange, targeting existing shortcoming of ICN-VANETs. Upcoming work, similarly to the process taken in this Thesis, can be divided into three areas of focus, namely, architecture, public transit and caching. The proposal of architecture-wide improvements such as distributed service tracking mechanisms can further increase the potential for service exchange. The discussion of public transit presented serves as the technical basis necessary for specialized research on the topic ot PTS, including definition and evaluation of specialized services and integration into currently existing application scenarios. The work presented is also a starting point for a larger discussion on the relevance of public transit, which we predict will only become more relevant as vehicular networks come to fruition. In regards to caching, research potential is predictably large, with research opportunities in improving cache coordination within vehicular networks. We observe potential not only in popularity and locality-based cache policy implementation but also in terms of service classification and other access prediction schemes As future work, the following topics are planned for investigation:

- One of the remaining challenges associated with service exchange is the storage and tracking of service provisioning records. By providing specialized mechanisms that can track and reward participation and collaboration, we expect the application scope of service exchange to be expanded. Towards this goal, design a privacy-aware blockchain-based service bartering models that take into account network usage and application classification. In tandem, we believe that time sensitive service management and predictive load management modeling can be applied as part of QoS targeted towards the different application classes. Lastly, as a means of supporting the knowledge and service schemes, service exchange monetization models are viewed as an important aspect that warrants additional research.

- Investigate request patterns and request-prediction at both packet and chunk-levels. Evaluation the effects of request aggregation on cache responses and definition of heuristics to improve communication with request batching;

- Investigate the benefits of additional request patterns and potential embedding of such knowledge in beacon messages. From this, design a predictive request batching algorithm to fulfill expected correlated requests;
- Design a mechanism that can autonomously detect location correlation in applications with non-explicit location properties. Establish metrics that determine location correlation aided by machine learning and big-data analysis tools.

Bibliography

[1] J. Saltzer, D. Reed, and D. Clark, "End-to-end arguments in system design," *ACM Transactions on Computer Systems (TOCS)*, vol. 2, pp. 277–288, November 1984.

[2] A. Boukerche, Sungbum Hong, and T. Jacob, "An efficient synchronization scheme of multimedia streams in wireless and mobile systems," *IEEE Transactions on Parallel and Distributed Systems*, vol. 13, pp. 911–923, September 2002.

[3] M. Handley, "Why the internet only just works," *BT Technology Journal*, vol. 24, pp. 119–129, July 2006.

[4] G. Xylomenos, C. N. Ververidis, V. A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, and G. C. Polyzos, "A survey of information-centric networking research," *IEEE Communications Surveys Tutorials*, vol. 16, pp. 1024–1049, July 2014.

[5] M. Conti, C. Boldrini, S. S. Kanhere, E. Mingozzi, E. Pagani, P. M. Ruiz, and M. Younis, "From MANET to people-centric networking: Milestones and open research challenges," *Computer Communications*, vol. 71, pp. 1–21, November 2015.

[6] M. Amadeo, C. Campolo, A. Molinaro, and G. Ruggeri, "Content-centric wireless networking: A survey," *Computer Networks*, vol. 72, pp. 1–13, October 2014.

[7] M. Gritter and D. R. Cheriton, "An architecture for content routing support in the internet," in *Proc. 3rd USENIX Symposium on Internet Technologies and Systems*, vol. 1, pp. 37–48, March 2001.

[8] V. Jacobson, "A new way to look at networking," *Google Tech Talk*, vol. 30, August 2006.

[9] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *Proceedings of the 5th International*

Conference on Emerging Networking Experiments and Technologies, pp. 1–12, December 2009.

[10] A. Boukerche and S. Rogers, "GPS query optimization in mobile and wireless networks," in *Proceedings. Sixth IEEE Symposium on Computers and Communications*, pp. 198–203, July 2001.

[11] R. W. Coutinho and A. Boukerche, "Information-centric intelligent vehicular networks: Challenges and guidelines," in *Proceedings of the 21st ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, MSWiM '18, pp. 7–8, ACM, October 2018.

[12] J. Wang, R. Wakikawa, and L. Zhang, "DMND: Collecting data from mobiles using named data," in *Vehicular Networking Conference (VNC), 2010 IEEE*, pp. 49–56, December 2010.

[13] CISCO, "VNI global fixed and mobile internet traffic forecasts (2016-2021)," June 2017.

[14] CISCO, "The zettabyte era: Trends and analysis," white paper, CISCO Public, June 2017.

[15] I. C. Society, "Ieee standard for wireless access in vehicular environments (WAVE) - networking services," *IEEE Std 1609.3-2010 (Revision of IEEE Std 1609.3-2007)*, pp. 1–144, December 2010.

[16] K. Abrougui, A. Boukerche, and R. W. N. Pazzi, "Design and evaluation of context-aware and location-based service discovery protocols for vehicular networks," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, pp. 717–735, September 2011.

[17] D. Grewe, M. Wagner, and H. Frey, "A domain-specific comparison of information-centric networking architectures for connected vehicles," *IEEE Communications Surveys & Tutorials*, March 2018.

[18] K. Pentikousis, B. Ohlman, D. Corujo, G. Boggia, G. Tyson, E. Davies, A. Molinaro, and S. Eum, "Information-centric networking: baseline scenarios," tech. rep., RFC Editor, March 2015.

[19] L. Zhang, D. Estrin, J. Burke, V. Jacobson, J. D. Thornton, D. K. Smetters, B. Zhang, G. Tsudik, D. Massey, C. Papadopoulos, *et al.*, "Named data networking (ndn) project," tech. rep., Xerox Palo Alto Research Center-PARC, October 2010.

[20] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," in *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pp. 181–192, August 2007.

[21] P. TalebiFard, V. C. Leung, M. Amadeo, C. Campolo, and A. Molinaro, "Information-centric networking for VANETs," in *Vehicular ad hoc Networks*, pp. 503–524, Springer, 2015.

[22] P. A. A. Gutiérrez, M. Zitterbart, Z. Boudjemil, M. Ghader, G. H. Garcia, M. Johnsson, A. Karouia, G. Lazar, M. Majanen, P. Mannersalo, D. Martin, M. T. Nguyen, S. P. Sanchez, P. Phelan, M. P. de Leon, G. Schultz, M. Söllner, Y. Zaki, and L. Zhao, "4WARD - architecture and design for the future internet," d-2.3.1, June 2010.

[23] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking," *IEEE Communications Magazine*, vol. 50, pp. 26–36, July 2012.

[24] S. H. Bouk, S. H. Ahmed, and D. Kim, "Vehicular content centric network (VCCN): a survey and research challenges," in *Proceedings of the 30th Annual ACM Symposium on Applied Computing*, pp. 695–700, ACM, 2015.

[25] Y. Zhang, A. Afanasyev, J. Burke, and L. Zhang, "A survey of mobility support in named data networking," in *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 83–88, April 2016.

[26] M. Amadeo, C. Campolo, and A. Molinaro, "Information-centric networking for connected vehicles: a survey and future perspectives," *IEEE Communications Magazine*, vol. 54, pp. 98–104, February 2016.

[27] M. A. Yaqub, S. H. Ahmed, S. H. Bouk, and D. Kim, "Interest forwarding in vehicular information centric networks: a survey," in *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, pp. 724–729, ACM, 2016.

[28] A. Ioannou and S. Weber, "A survey of caching policies and forwarding mechanisms in information-centric networking," *IEEE Communications Surveys & Tutorials*, vol. 18, pp. 2847–2886, May 2016.

[29] G. Zhang, Y. Li, and T. Lin, "Caching in information centric networking: A survey," *Computer Networks*, vol. 57, pp. 3128–3141, August 2013. Information Centric Networking.

[30] M. Zhang, H. Luo, and H. Zhang, "A survey of caching mechanisms in information-centric networking," *IEEE Communications Surveys Tutorials*, vol. 17, pp. 1473–1499, Third Quarter 2015.

[31] R. Tourani, T. Mick, S. Misra, and G. Panwar, "Security, privacy, and access control in information-centric networking: A survey," *arXiv preprint arXiv:1603.03409*, 2016.

[32] E. G. AbdAllah, H. S. Hassanein, and M. Zulkernine, "A survey of security attacks in information-centric networking," *IEEE Communications Surveys & Tutorials*, vol. 17, pp. 1441–1454, January 2015.

[33] X. Liu, Z. Li, P. Yang, and Y. Dong, "Information-centric mobile ad hoc networks and content routing: A survey," *Ad Hoc Networks*, vol. 58, pp. 255–268, April 2017.

[34] R. W. Coutinho, A. Boukerche, and A. A. Loureiro, "Design guidelines for information-centric connected and autonomous vehicles," *IEEE Communications Magazine*, vol. 56, pp. 85–91, October 2018.

[35] U. C. Service, *Automotive Resource Guide: A Reference for U.S. Exporters*. Export.gov, 3rd ed., June 2014.

[36] F. Bai and B. Krishnamachari, "Exploiting the wisdom of the crowd: localized, distributed information-centric VANETs," *IEEE Communications Magazine*, vol. 48, pp. 138–146, May 2010.

[37] I. T. Unionm, "IMT vision framework and overall objectives of the future development of IMT for 2020 and beyond," *ITU Working Document 5D/TEMP/548-E*, February 2015.

[38] J. Varman and K. Pickens, "Samsung and verizon announce first 5G customer trials set to begin in Q2 2017," February 2017.

[39] "Samsung, qualcomm and verizon showcase rapid progress towards 5G NR commercialization," 2018.

[40] Qualcomm, "Accelerating c-v2x commercialization." Online, November 2018.

[41] G. A. Association, "An assessment of LTE-V2X (PC5) and 802.11p direct communications technologies for improved road safety in the EU," tech. rep., 5G Automotive Association, 2017.

[42] E. T. S. Institute, "Intelligent transport systems (its); osi cross-layer topics; part 8: Interface between security entity and network and transport layer."

[43] I. C. Society, "Ieee standard for wireless access in vehicular environments (wave)– over-the-air electronic payment data exchange protocol for intelligent transportation systems (its)," *IEEE Std 1609.11-2010*, pp. 1–62, January 2011.

[44] L. Wang, R. Wakikawa, R. Kuntz, R. Vuyyuru, and L. Zhang, "Data naming in vehicle-to-vehicle communications," in *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on*, pp. 328–333, March 2012.

[45] S. Kumar, L. Shi, N. Ahmed, S. Gil, D. Katabi, and D. Rus, "Carspeak: A content-centric network for autonomous driving," *ACM SIGCOMM Computer Communication Review*, vol. 42, pp. 259–270, October 2012.

[46] Y. Yu, X. Li, M. Gerla, and M. Sanadidi, "Scalable VANET content routing using hierarchical bloom filters," in *2013 9th International Wireless Communications and Mobile Computing Conference (IWCMC)*, pp. 1629–1634, July 2013.

[47] D. Pesavento, G. Grassi, C. E. Palazzi, and G. Pau, "A naming scheme to represent geographic areas in NDN," in *Wireless Days (WD), 2013 IFIP*, pp. 1–3, November 2013.

[48] G. Grassi, D. Pesavento, G. Pau, R. Vuyyuru, R. Wakikawa, and L. Zhang, "VANET via named data networking," in *Computer Communications Workshops (INFOCOM WKSHPS), 2014 IEEE Conference on*, pp. 410–415, April 2014.

[49] S. H. Bouk, S. H. Ahmed, and D. Kim, "Hierarchical and hash based naming with compact trie name management scheme for vehicular content centric networks," *Computer Communications*, vol. 71, pp. 73–83, November 2015.

[50] W. Quan, C. Xu, J. Guan, H. Zhang, and L. Grieco, "Social cooperation for information-centric multimedia streaming in highway VANETs," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2014 IEEE 15th International Symposium on a*, pp. 1–6, June 2014.

[51] Z. Yan, S. Zeadally, and Y. J. Park, "A novel vehicular information network architecture based on named data networking (ndn)," *IEEE Internet of Things Journal*, vol. 1, pp. 525–532, December 2014.

[52] F. A. Silva, A. Boukerche, T. R. Silva, L. B. Ruiz, E. Cerqueira, and A. A. Loureiro, "Vehicular networks: A new challenge for content-delivery-based applications," *ACM Computing Surveys (CSUR)*, vol. 49, pp. 11:1–11:29, July 2016.

[53] F. Sailhan and V. Issarny, "Scalable service discovery for MANET," in *Third IEEE International Conference on Pervasive Computing and Communications*, pp. 235–244, March 2005.

[54] G. Grassi, D. Pesavento, L. Wang, G. Pau, R. Vuyyuru, R. Wakikawa, and L. Zhang, "Vehicular inter-networking via named data." Extended version of poster presented at HotMobile2013. Available on arXiv at: https://arxiv.org/pdf/1310.5980.pdf, October 2013.

[55] M. Amadeo, C. Campolo, and A. Molinaro, "Content-centric vehicular networking: An evaluation study," in *Network of the Future (NOF), 2012 Third International Conference on the*, pp. 1–5, November 2012.

[56] M. Amadeo, C. Campolo, and A. Molinaro, "CRoWN: content-centric networking in vehicular ad-hoc networks," *IEEE Communications Letters*, vol. 16, pp. 1380–1383, September 2012.

[57] G. Carofiglio, M. Gallo, and L. Muscariello, "ICP: Design and evaluation of an interest control protocol for content-centric networking," in *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on*, pp. 204–309, March 2012.

[58] M. Amadeo, C. Campolo, and A. Molinaro, "Design and analysis of a transport-level solution for content-centric VANETs," in *2013 IEEE International Conference on Communications Workshops (ICC)*, pp. 532–537, June 2013.

[59] M. Chen, D. O. Mau, Y. Zhang, T. Taleb, and V. C. Leung, "VENDNET: Vehicular named data network," *Vehicular Communications*, vol. 1, pp. 208–213, October 2014.

[60] S. H. Bouk, M. A. Yaqub, S. H. Ahmed, and D. Kim, "Evaluating interest/data propagation in vehicular named data networks," in *Proceedings of the 2015 Conference on Research in Adaptive and Convergent Systems*, pp. 256–259, October 2015.

[61] S. H. Ahmed, S. H. Bouk, and D. Kim, "RUFS: Robust forwarder selection in vehicular content-centric networks," *IEEE Communications Letters*, vol. 9, pp. 1616–1619, September 2015.

[62] C. Bian, T. Zhao, X. Li, and W. Yan, "Boosting named data networking for data dissemination in urban VANET scenarios," *Vehicular Communications*, vol. 2, pp. 195–207, October 2015.

[63] Y. Yu, Y. Li, X. Ma, W. Shang, M. Sanadidi, and M. Gerla, "Scalable opportunistic VANET content routing with encounter information," in *2013 21st IEEE International Conference on Network Protocols (ICNP)*, pp. 1–6, October 2013.

[64] X. Liu, M. João Nicolau, A. Costa, J. Macedo, and A. Santos, "A geographic opportunistic forwarding strategy for vehicular named data networking," in *Intelligent Distributed Computing IX: Proceedings of the 9th International Symposium on Intelligent Distributed Computing – IDC'2015, Guimarães, Portugal, October 2015*, vol. 616, pp. 509–521, Springer International Publishing, 2016.

[65] X. Yu, R. W. Coutinho, A. Boukerche, and A. A. Loureirol, "A distance-based interest forwarding protocol for vehicular information-centric networks," in *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, pp. 1–5, IEEE, October 2017.

[66] R. W. Coutinho, A. Boukerche, and X. Yu, "A novel location-based content distribution protocol for vehicular named-data networks," in *2018 IEEE Symposium on Computers and Communications (ISCC)*, pp. 01007–01012, IEEE, June 2018.

[67] G. Grassi, D. Pesavento, G. Pau, L. Zhang, and S. Fdida, "Navigo: Interest forwarding by geolocations in vehicular named data networking," in *World of Wire-*

less, Mobile and Multimedia Networks (WoWMoM), 2015 IEEE 16th International Symposium on a, pp. 1–10, June 2015.

[68] N. Laoutaris, S. Syntila, and I. Stavrakakis, "Meta algorithms for hierarchical web caches," in *Performance, Computing, and Communications, 2004 IEEE International Conference on*, pp. 445–452, 2004.

[69] J. Xu, X. Tang, and D. L. Lee, "Performance analysis of location-dependent cache invalidation schemes for mobile environments," *IEEE Transactions on Knowledge and Data Engineering*, vol. 15, pp. 474–488, March 2003.

[70] W. K. Chai, D. He, I. Psaras, and G. Pavlou, "Cache less for more in information-centric networks," in *International Conference on Research in Networking*, May 2012.

[71] S. Eum, K. Nakauchi, M. Murata, Y. Shoji, and N. Nishinaga, "CATT: Potential based routing with content caching for ICN," in *Proceedings of the Second Edition of the ICN Workshop on Information-centric Networking*, ICN '12, (New York, NY, USA), pp. 49–54, ACM, 2012.

[72] I. Psaras, W. K. Chai, and G. Pavlou, "Probabilistic in-network caching for information-centric networks," in *Proceedings of the Second Edition of the ICN Workshop on Information-centric Networking*, ICN '12, pp. 55–60, ACM, February 2012.

[73] F. Kuliesius and P. Paulauskas, "Simulation of content caching in information centric networking," in *2015 Seventh International Conference on Ubiquitous and Future Networks*, pp. 391–394, July 2015.

[74] Y. Kim and I. Yeom, "Performance analysis of in-network caching for content-centric networking," *Computer Networks*, vol. 57, pp. 2465–2482, September 2013.

[75] Y.-T. Yu and M. Gerla, "Potential benefits of information-centric networks for VANETs." ResearchgGate, February 2015.

[76] M. Abuelela, S. Olariu, and M. Weigle, "NOTICE: An architecture for the notification of traffic incidents," in *Vehicular Technology Conference, 2008. VTC Spring 2008. IEEE*, pp. 3001–3005, May 2008.

[77] I. C. Society, "IEEE standard for information technology– local and metropolitan area networks– specific requirements– part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications amendment 6: Wireless access in vehicular environments," *IEEE Std 802.11p-2010 (Amendment to IEEE Std 802.11-2007 as amended by IEEE Std 802.11k-2008, IEEE Std 802.11r-2008, IEEE Std 802.11y-2008, IEEE Std 802.11n-2009, and IEEE Std 802.11w-2009)*, pp. 1–51, July 2010.

[78] G. Pocovi, M. Lauridsen, B. Soret, K. I. Pedersen, and P. Mogensen, "Automation for on-road vehicles: Use cases and requirements for radio design," in *Vehicular Technology Conference (VTC Fall), 2015 IEEE 82nd*, pp. 1–5, IEEE, September 2015.

[79] A. Boukerche, S. Hong, and T. Jacob, "A distributed algorithm for dynamic channel allocation," *Mobile Networks and Applications*, vol. 7, pp. 115–126, April 2002.

[80] P. K. Sharma, S. Y. Moon, and J. H. Park, "Block-VN: A distributed blockchain based vehicular network architecture in smart city," *Journal of Information Processing Systems*, vol. 13, pp. 184–195, March 2017.

[81] J. B. Kenney, "Dedicated short-range communications (DSRC) standards in the united states," *Proceedings of the IEEE*, vol. 99, pp. 1162–1182, July 2011.

[82] Z. Xu, X. Li, X. Zhao, M. H. Zhang, and Z. Wang, "DSRC versus 4G-LTE for connected vehicle applications: a study on field experiments of vehicular communication performance," *Journal of Advanced Transportation*, vol. 2017, 2017.

[83] M. Amadeo, C. Campolo, and A. Molinaro, "Enhancing content-centric networking for vehicular environments," *Computer Networks*, vol. 57, pp. 3222–3234, November 2013.

[84] M. O'Rielly, "Defining auto safety of life in 5.9GHz." Available At: `https://www.fcc.gov/news-events/blog/2016/06/08/defining-auto-safety-life-59-ghz`, June 2016.

[85] A. Boukerche and C. Tropper, "A distributed graph algorithm for the detection of local cycles and knots," *IEEE Transactions on Parallel and Distributed Systems*, vol. 9, pp. 748–757, August 1998.

[86] P. TalebiFard, R. Ravindran, A. Chakraborti, J. Pan, A. Mercian, G. Wang, and V. Leung, "An information centric networking approach towards contextualized edge service," in *Consumer Communications and Networking Conference (CCNC), 2015 12th Annual IEEE*, pp. 250–255, January 2015.

[87] C. Rezende, A. Boukerche, H. S. Ramos, and A. A. Loureiro, "A reactive and scalable unicast solution for video streaming over VANETs," *IEEE Transactions on Computers*, vol. 64, pp. 614–626, March 2015.

[88] E. Lee, E.Lee, M. Gerla, and S. Oh, "Vehicular cloud networking: Architecture and design principles," *IEEE Communications Magazine*, vol. 52, pp. 148–155, February 2014.

[89] J. Chen, S. Li, H. Yu, Y. Zhang, D. Raychaudhuri, R. Ravindran, H. Gao, L. Dong, G. Wang, and H. Liu, "Exploiting ICN for realizing service-oriented communication in IoT," *IEEE Communications Magazine*, vol. 64, pp. 20–30, December 2016.

[90] M. Amadeo, C. Campolo, and A. Molinaro, "Named data networking for priority-based content dissemination in VANETs," in *Personal, Indoor, and Mobile Radio Communications (PIMRC), 2016 IEEE 27th Annual International Symposium on*, pp. 1–6, December 2016.

[91] W. He, G. Yan, and L. D. Xu, "Developing vehicular data cloud services in the iot environment," *IEEE Transactions on Industrial Informatics*, vol. 10, pp. 1587–1595, May 2014.

[92] N. Cordeschi, D. Amendola, and E. Baccarelli, "Reliable adaptive resource management for cognitive cloud vehicular networks," *IEEE Transactions on Vehicular Technology*, vol. 64, pp. 2528–2537, June 2015.

[93] P. Talebifard and V. Leung, "Towards a content-centric approach to crowd-sensing in vehicular clouds," *ournal of Systems Architecture*, vol. 59, pp. 976–984, November 2013.

[94] X. Hu, L. Wang, Z. Sheng, P. TalebiFard, L. Zhou, J. Liu, and V. Leung, "Towards a service centric contextualized vehicular cloud," in *Proceedings of the Fourth ACM International Symposium on Development and Analysis of Intelligent Vehicular Networks and Applications*, DIVAnet '14, pp. 73–80, November 2014.

[95] A. Boukerche, C. Rezende, and R. W. Pazzi, "Improving neighbor localization in vehicular ad hoc networks to avoid overhead from periodic messages," in *GLOBECOM 2009 - 2009 IEEE Global Telecommunications Conference*, pp. 1–6, November 2009.

[96] L. A. Villas, A. Boukerche, D. L. Guidoni, H. A. De Oliveira, R. B. De Araujo, and A. A. Loureiro, "An energy-aware spatio-temporal correlation mechanism to perform efficient data collection in wireless sensor networks," *Computer Communications*, vol. 36, pp. 1054–1066, May 2013.

[97] Y. Du, L. Zhang, Y. Feng, Z. Ren, and Z. Wang, "Performance analysis and enhancement of IEEE 802.11p/1609 protocol family in vehicular environments," in *13th International IEEE Conference on Intelligent Transportation Systems*, pp. 1085–1090, September 2010.

[98] M. Slavik, I. Mahgoub, and M. M. Alwakeel, "Analysis of beaconing message rate in VANET multi-hop broadcast protocols," in *High Capacity Optical Networks and Emerging/Enabling Technologies*, pp. 37–41, December 2012.

[99] R. Oliveira, C. Montez, A. Boukerche, and M. Wangham, "Reliable data dissemination protocol for VANET traffic safety applications," *Ad Hoc Networks*, vol. 63, pp. 30–44, August 2017.

[100] A. Vargas, "The OMNeT++ discrete event simulation system," in *Proceedings of the European simulation multiconference (ESM'01)*, vol. 9, p. 65, June 2001.

[101] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, vol. 10, pp. 3–15, January 2011.

[102] D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker, "Recent development and applications of SUMO - Simulation of Urban MObility," *International Journal On Advances in Systems and Measurements*, vol. 5, December 2012.

[103] The City of Edmonton, "Traffic flow maps & reports.." Available At: `http://www.edmonton.ca/transportation/traffic_reports/traffic-reports-flow-maps.aspx`, 2010–2015.

[104] City of Calgary, "Historic traffic volume flow map – AAWT maps." Available At: `http://www.calgary.ca/Transportation/TP/Pages/Planning/Transportation-Data/Traffic-volume-flow-maps.aspx`, 1963–2017.

[105] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: Evidence and implications," in *INFOCOM'99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1, pp. 126–134, IEEE, March 1999.

[106] G. Yan, S. Olariu, and M. C. Weigle, "Providing VANET security through active position detection," *Computer Communications*, vol. 31, pp. 2883–2897, July 2008.

[107] R. G. Engoulou, M. Bellaïche, S. Pierre, and A. Quintero, "VANET security surveys," *Computer Communications*, vol. 44, pp. 1–13, May 2014.

[108] C. Dannewitz, "Netinf: An information-centric design for the future internet," in *Proc. 3rd GI/ITG KuVS Workshop on The Future Internet*, May 2009.

[109] H. Ko, Y. Kim, D. Suh, and S. Pack, "A proactive content pushing scheme for provider mobility support in information centric networks," in *2014 IEEE 11th Consumer Communications and Networking Conference (CCNC)*, pp. 523–524, January 2014.

[110] T. Wei, L. Chang, B. Yu, and J. Pan, "MPCS: A mobility/popularity-based caching strategy for information-centric networks," in *2014 IEEE Global Communications Conference*, pp. 4629–4634, December 2014.

[111] K. Cho, M. Lee, K. Park, T. T. Kwon, Y. Choi, and S. Pack, "WAVE: Popularity-based and collaborative in-network caching for content-oriented networks," in *2012 Proceedings IEEE INFOCOM Workshops*, pp. 316–321, March 2012.

[112] J. Iqbal and P. Giaccone, "Interest-based cooperative caching in multi-hop wireless networks," in *Globecom Workshops (GC Wkshps), 2013 IEEE*, IEEE, December 2013.

[113] H. Jin, D. Xu, C. Zhao, and D. Liang, "Information-centric mobile caching network frameworks and caching optimization: A survey," *EURASIP Journal on Wireless Communications and Networking*, vol. 2017, p. 33, February 2017.

[114] D. L. Lee, J. Xu, B. Zheng, and W. C. Lee, "Data management in location-dependent information services," *IEEE Pervasive Computing*, vol. 1, pp. 65–72, July 2002.

[115] M. Wang, A. Ailamaki, and C. Faloutsos, "Capturing the spatio-temporal behavior of real traffic data," *Performance Evaluation*, vol. 49, pp. 147–163, September 2002.

[116] N. Sluijs, F. Iterbeke, T. Wauters, F. De Turck, B. Dhoedt, and P. Demeester, "Cooperative caching versus proactive replication for location dependent request patterns," in *Journal of Network and Computer Applications*, vol. 34, pp. 562–574, Elsevier, March 2011.

[117] K. Machado, A. Boukerche, E. Cerqueira, and A. A. F. Loureiro, "Long-term spatiotemporal analysis of social media for device-to-device networks," in *2016 IEEE Global Communications Conference (GLOBECOM)*, pp. 1–6, December 2016.