University of
Nottingham
UK | CHINA | MALAYSIA

Resource-Aware Trust-Based Security for Vehicular Smart-Grid Networks

Thesis submitted to the University of Nottingham for the degree of

Doctor of Philosophy, September 2023

Adam David Walker

Supervised by

Milena Radenkovic

Chris Greenhalgh

Abstract

Vehicular smart grids (VSGs) are localised electrical grids that form spontaneously when heterogeneous nodes, including electric vehicles (EVs) and charging stations (CSs), are temporarily co-located and can communicate and exchange energy bi-directionally, e.g., vehicle-to-vehicle (V2V). VSGs are highly dynamic due primarily to EV mobility. Instability and lack of trust amongst nodes make energy management challenging and create opportunities for malicious actors to disrupt supply through energy denial-of-service (EDoS) attacks. We see VSGs as an extension of opportunistic networks (OppNets). This thesis proposes CognitiveCharge, a framework and template protocol for independent, mutually untrusted nodes to coordinate localised, opportunistic energy exchange that builds on data routing strategies in OppNets. Each device in our VSG model operates as an independent CognitiveCharge node using real-time utility-driven decision-making and cross-layer predictive analytics using first and second-hand observations. We implement CognitiveCharge by significantly extending existing agent-based discrete event network simulation software. CognitiveCharge performance is explored across a range of multi-day urban and semi-urban VSG scenarios, which include real-world and pseudo-realistic data and were developed specifically for this work. Our simulation-based experiments show that CognitiveCharge increases the availability of energy for EVs to expend on mobility, even when under an active EDoS attack. CognitiveCharge nodes can identify and exploit energy exchange opportunities to increase local and regional availability of on-demand energy as well as mitigate the impact of EDoS attacks in terms of energy loss by accurately detecting and avoiding exchanges with malicious nodes.

Abbreviations

CD	Contact Duration
CSS	Cyber-Physical System
CS	Charging Station
DoS	Denial-of-Service
DTN	Delay -Tolerant Network
EDoS	Energy Denial-of-Service
EV	Electric Vehicle
FIFO	First-In-First-Out
GPS	Global Positioning System
G2V	Grid-to-Vehicle
ICEV	Internal Combustion Engine Vehicle
ICT	Inter-Contact Time
IoE	Internet of Energy
IoT	Internet of Things
LIFO	Last-In First-Out
OppNet	Opportunistic Network
P2P	Peer-to-Peer

POI	Point of Interest
RER	Renewable Energy Resource
SG	Smart Grid
UAV	Unmanned Aerial Vehicle
UUV	Unmanned Underwater Vehicle
V2G	Vehicle-to-Grid
V2V	Vehicle-to-Vehicle
V2X	Vehicle-to-Anything
VSG	Vehicle Smart Grid

Acknowledgements

I would like to express my profound gratitude and heartfelt appreciation to my dedicated supervisors, Milena Radenkovic and Chris Greenhalgh. Your unwavering support, invaluable guidance, and consistent encouragement have been fundamental to the successful completion of this work. I am extraordinarily grateful to you both for your patience, kindness, and expertise throughout my journey.

I also extend my thanks to Bob John. The insightful discussions we had during the early phase of this project helped shape the direction and focus of this research.

Finally, I am thankful for the support and motivation of my friends and colleagues at the University of Nottingham.

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Reference 2049961.

Table of Contents

Chapter 1

Introduction

1.1 Background and Motivation

The envisioned smart grid (SG) paradigm promises to revolutionise the existing electrical grid infrastructure and energy exchange marketplaces. As the world faces significant challenges from climate change whilst facing an ever-increasing demand for energy, SG technologies seek to provide innovative solutions to address a wide range of concerns and harmonise supply and demand dynamics. From increased localisation of power generation and supply to greater penetration of green and renewable energy, the much-lauded SG paradigm is positioned to be one of the most important, large-scale, global developments of the 21st century. As such, the SG has garnered significant attention and remains the focus of many researchers, industrial interests, and governments around the world. Despite this, there remains considerable work to be done.

The SG is facilitated through the deep integration of highly heterogeneous devices via data communications and localised energy exchange mechanisms. Through these, SG nodes collaboratively coordinate and manage energy flows amongst many nodes, such as IoT devices, EVs, and local RERs, and stakeholders, such as device owners and utility companies. The benefits of the SG are widely documented in the literature (e.g. efficiency, sustainability, stability, green energy), but there are also significant barriers in order to fully realise it.

The traditional electrical grid is strictly hierarchical. Industrial, commercial, and domestic consumers acquire electrical power from upstream regional distributors, who, in turn, obtain electricity from conventional generation facilities. This rigid production-to-distribution-to-consumption supply chain means the stakeholders involved have fixed roles in the system. Power plants are responsible for generating the energy required by the grid. Large distribution networks handle the transmission of electrical power to regional distribution companies and are responsible for overall grid stability. Local energy distribution companies are accountable for supplying electricity to customers, providing metering and billing. Conversely, the smart grid paradigm embraces dynamic and flexible roles, supported via bidirectional flows of energy and information amongst highly heterogeneous devices, ranging from smart home appliances, and mobile devices to autonomous EVs and local renewable energy resources.

The particular aspects of the SG that this work considers are well represented by the 'internet of energy' (IoE) term, which describes the deep integration of smart grid technologies with such devices and aims to facilitate efficient, real-time, adaptive coordination of energy supply and demand in localised SGs. Networked communications and intelligent technologies at the SG edge enable energy to be moved on demand to meet continually changing requirements within the SG system. For example, a residential consumer with solar panels may act as a temporary prosumer at times when they produce more energy locally than they consume. At these times, the prosumer can choose to sell any surplus energy back to the grid, allowing them to reduce their overall energy expenditure whilst the wider grid simultaneously benefits from increased stability and local energy availability.

As to the ability of devices and appliances, such as smart thermostats and electric vehicles, to communicate with the electrical grid in real time, enabling better management of energy supply and demand. This integration allows for more efficient distribution of energy, increased use of renewable energy, and improved grid stability. The internet of energy is an important component of the smart grid paradigm and is a rapidly developing area of research and innovation [1].

Figure 1 depicts a high-level overview of the traditional and SG paradigms. VSGs sit at the edge of the future SG and are formed primarily of EVs and infrastructure CSs, such as those at homes, car parks, and dedicated recharging stations. As illustrated in Figure 1, heterogeneous VSGs form spontaneously amongst temporarily collocated nodes via localised communications. Due to vehicular mobility, the VSG is highly dynamic. Increased availability of convenient charging infrastructure is important to consumer adoption of EVs [2], [3] and the VSG presents increased opportunities to acquire and offload energy amongst EVs and with the wider grid. The VSG paradigm also offers the potential for increased grid stability and adaptability to the challenging dynamics of local and regional energy flux [4], [5]. Specific technologies for V2V energy exchange are beyond the scope of this work, but there is a vast variety of close-range energy exchange proposals in the literature, including wired and wireless mechanisms [6]. Rather than being dependent on a particular technology, we only consider that nodes must be physically adjacent in order to exchange energy.

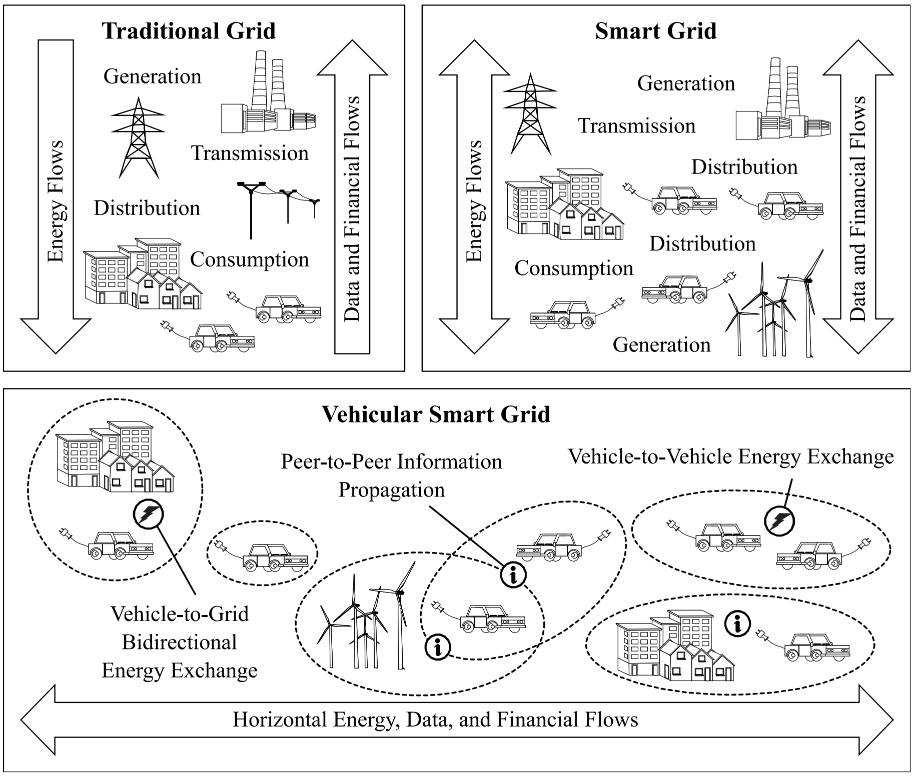

Figure 1 Electrical Grid Paradigms

Complex-adaptive cyber-physical systems (CPSs) such as the future SG, have an increased attack surface than their traditional counterparts as widely networked control processes expose new and previously unexploitable vulnerabilities [7], [8], [9]. Compounding this, the existing electrical grid infrastructure is already facing significant stability and security challenges which both increase susceptibility and reduce resilience to attack [7], [10]. Meanwhile, ongoing energy theft from the electrical grid remains a substantial problem in both developed and developing nations, with global losses of approximately 25 billion dollars and rates of theft reaching as high as 50 per cent in some markets [11], [12].

In recent years, security researchers have observed rising numbers of sophisticated attacks against the energy sector conducted by a wide range of groups who are increasingly targeting more localised grid operations [8]. There are numerous well-documented cases where cyberattacks have targeted energy infrastructure. Several large-scale coordinated cyberattacks against energy distribution networks exploited networked control systems to disrupt power to almost a quarter of a million customers, impacting households, industry, and emergency services [13], [14]. The high capacity and energy demands of EVs make them a valuable target for attackers and their integration into the SG has opened up new avenues for malicious actors. A range of novel attacks targeting EV charging systems has been demonstrated to have a severe impact on the stability of electrical power distribution [15].

Existing approaches to energy exchange in vehicular smart grids are reliant on a high degree of connectivity, with fully or semi-centralised completely trusted infrastructure is necessary to provide security assurances and coordination of distributed nodes. Furthermore, many existing approaches are reliant on coarse-grained estimations of energy usage combined with limited advanced knowledge for scheduling. We argue that centralised decision-making for V2V and V2G charging does not perform well in highly dynamic, distributed, grid-edge scenarios due to limited scalability, responsiveness, real-time adaptability, and fairness. Research has shown that centralised optimisation and global optimum approaches are unsuited to environments with highly dynamic topologies, fluctuating geo-temporal energy availability, and frequent temporal disconnections – such as vehicular smart grid networks – due to such the assumption of a priori knowledge and algorithms which often unfairly disadvantage some nodes [16], [17]. Locally and centrally optimised algorithms can be outperformed by collaborative approaches in networks with complex dynamic spatiotemporal topologies [17].

It is the aim of this thesis to design and validate a framework for facilitating P2P energy exchange amongst nodes in untrusted heterogeneous smart vehicular grid environments and in the presence of an active energy depletion attack. Such a framework is a necessary prerequisite to realising fully opportunistic energy exchange which is adaptive to the manifold inherent real-time socio-spatiotemporal dynamics. Meeting the challenges of providing security of energy and communications exchange in the presence of both malicious and non-malicious threats is recognised as being fundamental to the success of the future smart grid [18]. Furthermore, distributed, networked security solutions are identified as crucial to realising a responsive and self-healing smart grid however fully implementing such systems raises major technological challenges [19], [20].

This section has provided background and motivation for the research, introducing some aspects of P2P energy exchange and energy depletion attacks in VSGs. The remainder of this chapter is arranged as follows. Firstly we detail the vehicular smart grid environment central to this thesis; here we also specify the core scenario presumptions and delineates the scope of this work. Expanding on this, we describe the specifics of the energy depletion attack considered in this thesis. The precise research focus is then given, followed by the main contributions of this work. Finally, an overview of the structure of this thesis is provided.

1.2 Vehicular Smart Grid

In this work we consider a heterogeneous vehicular smart grid (VSG) environment comprising a diverse range of electric vehicles (EVs) and electrical infrastructure charging stations (CSs). The set of participating nodes in the vehicular SG consists primarily of high energy capacity mobile EVs and static electrical grid infrastructure access points, for example, at homes and charging stations. There is inherently a large degree of real-time variable heterogeneity across the physical attributes of nodes in VSG scenarios. EVs, such as consumer cars, motorcycles, and commercial goods vehicles, vary according to computational capability, energy storage capacity, characteristics of available data communications and energy exchange interfaces, rates of energy consumption for service provision, and dynamic patterns of mobility. In addition to the intrinsic heterogeneity at the physical and topological layers, we also consider each node to be fully independent, having its own dynamic internal motivations which govern its energy resource related behaviour.

Unlike in scenarios with externally operated nodes, e.g., cooperative managed fleet vehicles, in our heterogeneous vehicular smart grid scenario we presume that objectives of nodes are not shared and may even be in direct conflict. For example, consider a simplified model of an autonomous electric taxicab. To provide financially viable carriage services, it must set prices for journeys which are competitive with rivals whilst also being adequately offset against operating expenses. In order to accomplish this, the taxicab will continually adapt its behaviour, maximising revenue from service provision whilst simultaneously minimising the cost of energy acquisition in terms of both the raw energy price and the loss of revenue from downtime whilst charging. As a result of this, it is likely that the taxicab will seek to prioritise its acquisition of energy from low-cost suppliers during quieter operating periods, only using higher cost, slow supply sources when it deems necessary. Contrast the taxicab example with an emergency service vehicle model, whose primary purpose is to get to locations as fast as possible and with minimal downtime. Such a vehicle will almost exclusively favour nearby fast charging energy suppliers, disregarding any associated increased costs. Our VSG scenario assumes the inclusion of many kinds of vehicles and CSs, permitting the participation of any node seeking to acquire or offload energy.

We broadly summarise asynchronous opportunistic energy exchange in the VSG as a three-phase, supplier-driven, sequential process. The work in this thesis is intended to complement existing approaches and easily integrate with active and future technologies. In order to remain broadly compatible, only essential presumptions regarding the scenario are made. As illustrated in Figure 2, these steps are as follows:

- **Initialisation and Negotiation**: A node with surplus energy to offload (a) advertises itself as an energy supplier to collocated neighbours (b and c), who can then respond by sending requests for energy to the supplier. The supplier selects a suitable seeker from those interested – providing one exists.
- **P2P Energy Exchange**: Once the supplier has selected a seeker (in this case, node b), energy is supplied in a P2P manner, typically in exchange for financial return. This can happen either directly (such as using wireless or wired energy exchange between the two nodes) or via auxiliary mechanisms.
- **Peer Exchange Acknowledgement**: After the exchange, both the energy supplier (a) and consumer (b) validate the transaction by propagating records to the wider network (c). This verifies the transaction and the newly increased balance of the supplier's available funds.

The initialisation phase may include some negotiation of terms between parties regarding the details of the exchange, for example, the precise amount of energy to transfer and the associated price of energy. This thesis does not propose a specific approach for negotiation, nor does the work in this thesis restrict negotiation in advance of, or during, the exchange of energy between peers. Similarly, the physical exchange of energy between hosts can be direct or utilise available auxiliary hardware as an intermediary. Various approaches to localised energy exchange amongst EVs and grid infrastructure are being actively explored and deployed. The work in this thesis is purposefully agnostic to the specifics of the P2P energy exchange and therefore supports heterogeneous exchange mechanisms. We presume only that collocated nodes can exchange energy and data locally in real-time via some compatible means.

Finally, in line with existing trends for increasing privacy, we presume a privacy-aware scenario in which protocols only expose the minimally necessary amount of data to nodes via monitoring of exchange information and data propagation. Table 1 details the amount of information regarding an energy exchange event that we consider could visible to VSG nodes, based on their proximity. Participants directly involved in the energy exchange (viz. those sending or receiving energy) have full awareness of the physical identity of their peers, the location of the exchange, the precise amount of energy that was transferred, and the data and metadata communicated about exchanged, such as the price of the energy purchased. This information is necessary for both parties to exchange energy in a P2P manner within the VSG.

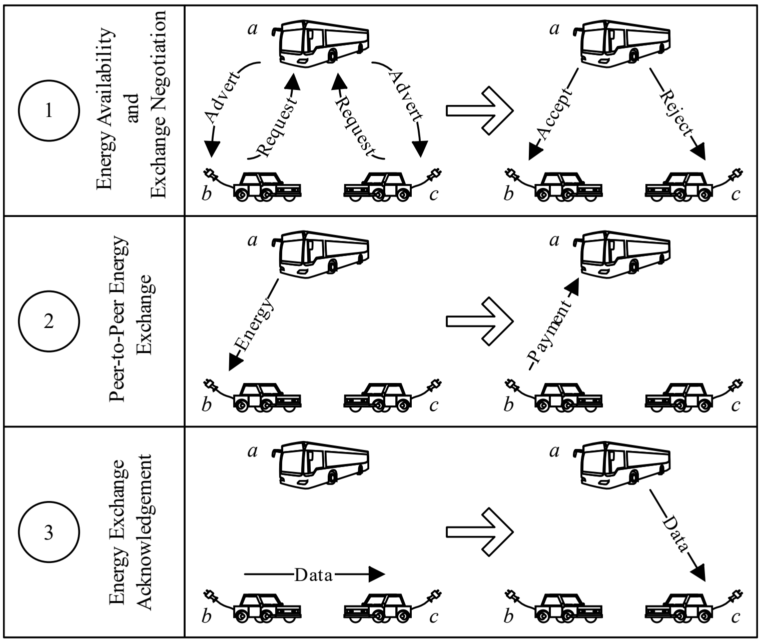

1
Energy Availability and Exchange Negotiation
2
Peer-to-Peer Energy Exchange
3
Energy Exchange Acknowledgement
a
b
c
Advert
Request
Request
Advert
Accept
Reject
Energy
Payment
Data
Data

Figure 2 VSG Energy Exchange Process

The collocated first-hand observers of the exchange in Table 1 have awareness of the nodes participating in the exchange, the role of each as sender or receiver, and the location. This information may not be readily published by the nodes participating in the exchange but can be inferred through other means. For instance, location can be physically observed and estimates of the amount of energy exchanged can be obtained through the time the nodes spent together. Confidantes – peers who receive explicitly propagated metadata concerning the transaction – are not sent information regarding the specific details of the exchange. Finally, for other nodes in the VSG we presume that without being informed of the exchange, these participants have no information explicitly shared or leaked to them.

The level of privacy of a given energy exchange protocol for the VSG can be laxer than the maximally stringent presumptions highlighted in Table 1. In a real-world scenario it could be possible for nodes to obtain information by a range of means, including aggregating observations, eavesdropping, data analysis etc. Nevertheless, we presume that additional data is not available and therefore cannot be relied upon for informing immediate and future decision making. This model increases the flexibility of our proposed CognitiveCharge framework and considers a highly privacy-oriented VSG scenario which is a core consideration of the future SG. By assuming strict, core constraints, we facilitate future works with more lax constraints to build more readily atop CognitiveCharge.

Table 1 VSG Energy Exchange Information Visibility

Information	**Participants**	**Observers**	**Confidantes**	**Others**
Identity	Yes	Yes	Yes	No
Location	Yes	Yes	No	No
Energy	Yes	No	No	No
Data	Yes	No	No	No

1.3 Energy Depletion Attack

Of concern for existing and future smart-grids is the potential for energy depletion attacks which, in addition to the immediate loss of power, can lead to financial loss and regional denial-of-service as communities become isolated from energy suppliers. In the vehicular smart grid, such an attack also represents a type of energy theft analogous to cheque fraud, wherein a malicious node takes advantage of the lack of global knowledge and regional network islanding to acquire energy at no cost. An energy denial-of-service (EDoS) attack considers nodes either maliciously, or through fault mechanisms, preventing nodes from accessing energy. Whilst this could be through blocking access, we consider the more devastating definition explored in the literature wherein a node continually acquires energy from a host until it is depleted (e.g. [21]).

Due to the high degree of heterogeneity, EDoS attacks can have a severe impact on the VSG. The safe operating margins of the electrical grid narrow year on year owing to rising demand and consequently the sensitivity to attack rises. In addition to the immediate loss of power which causes localised EDoS for the node, the impact of the attack expands beyond the immediate loss of energy for the non-malicious host. This can escalate to result in significant energy and financial losses and result in regional DoS as communities become isolated from energy suppliers.

We consider a malicious or hacked energy-seeking node that seeks to cause localised depletion of energy targeting an individual node, a geographic region, or a community of nodes. Absent of a priori knowledge or the ability to access a trusted authority, energy seekers and suppliers are vulnerable to a malicious node using falsified information to illegitimately obtain access to energy resources. As shown in Figure 3, in the first step, the malicious node presents itself as being the node most in need of energy to the advertising supplier, posing as a legitimate consumer seeking energy so as to position itself as the most desirable recipient (ahead of its peers). The second step then sees the attacking node acquire as much energy as possible from the supplier. As payments cannot be verified immediately, the attacker can make a false payment for the energy received in order to deceive the supplier. Finally, the malicious node disseminates false information which contradicts the energy supplier and seeks to invalidate any records of the exchange. This is possible because the network topology, which is inherently dynamic and disconnection prone, prohibits nodes from readily reaching consensus. The attack can be further exacerbated through attackers coordinating to suppress the occurrence of the exchange. As a result of the energy depletion attack, nodes which were dependent upon the supplier for energy resources are unable to provide services due to lack of availability.

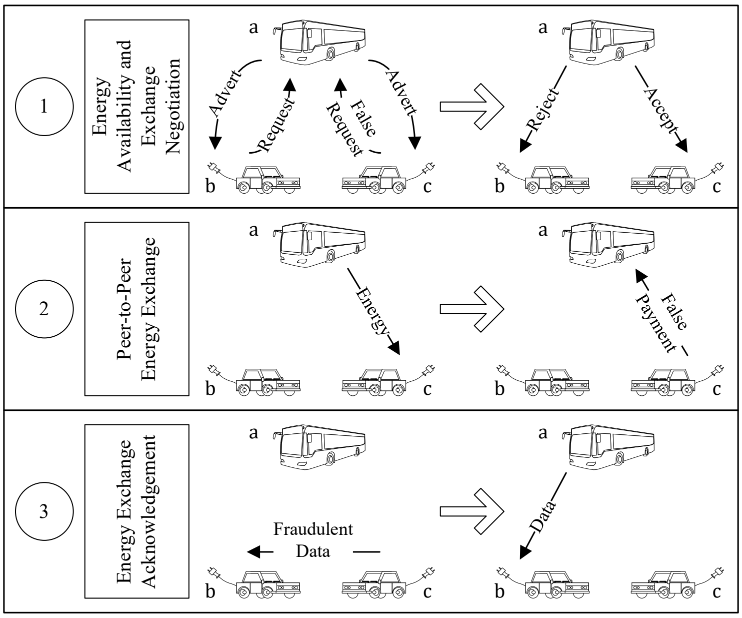
1
Energy Availability and Exchange Negotiation
a
Advert
Request
False Request
Advert
b
c
a
Reject
Accept
b
c
2
Peer-to-Peer Energy Exchange
a
Energy
b
c
a
False Payment
b
c
3
Energy Exchange Acknowledgement
a
Fraudulent Data
b
c
a
Data
b
c

Figure 3 Energy Depletion Attack

1.4 Research Focus

This thesis is concerned with adaptively balancing the competing trade-offs between energy and security such that VSG nodes can continue to share energy in untrusted environments and whilst under active attack by malicious nodes. The overarching aim of this work revolves around exploring modelling energy analogous to data. More specifically, we aim to apply and build upon data routing decision-making principles from OppNets and DTNs to the VSG, as defined earlier in this chapter. Whilst recognising the differences, we seek to exploit the parallels between data and energy and aim to facilitate independent VSG nodes to make energy and security-aware decisions in untrusted VSGs to increase utility service provision. The precise focus of this thesis is as follows:

In heterogenous VSGs, with fully localised communication and decision-making, is it possible to increase the service provision utility of nodes and limit energy losses in the presence of an EDoS attack conducted by malicious nodes?

In this work, we consider VSG scenarios comprising mobile, roaming EVs and static CSs. For EVs, service provision utility, therefore, refers to the ability of nodes to meet the demands and expectations of their owners with respect to completing journeys. As highlighted earlier in this Chapter, this work focuses on the fully localised communication and decision-making that we have highlighted as being necessary to facilitate real-time energy exchange amongst mutually untrusted nodes in the VSG model central to this work. Under this model, nodes must independently make energy and security-aware decisions with respect to whether to exchange energy and with whom when presented with opportunities to do so. For example, consider two EVs that have no previous direct encounters. Upon meeting, they must each individually determine whether to do nothing, offer energy for exchange, or request a transfer of energy. The core research question can be rephrased and posited in the context of a given VSG scenario. For any given EV in the VSG, does it have sufficient battery capacity, or can it acquire sufficient energy from a peer to move to its next destination when needed, even when malicious nodes are actively conducting an EDoS attack? This encapsulates EVs (and CSs) also offloading energy when in surplus.

In addressing the overarching research question, we also consider the following sub-questions for the heterogeneous VSG scenarios considered in this work (comprising EVs and CSs):

- *How does a fully localised communication and decision-making-based approach to energy exchange affect the level of energy of EVs in heterogeneous VSG scenarios?*
- *What are the energy losses incurred by EVs under an EDoS attack and to what extent can this these losses be reduced by an attack detection and trust scheme?*
- *How do independent and collaborative trust management schemes impact the opportunities for energy exchange and availability of energy of EVs?*

In meeting these research questions and the overarching aims, we consider the following objectives: to develop a novel framework for peer energy exchange in the VSG; to devise a prototype protocol to facilitate this; to model a range of VSGs in suitable simulation software, including an EDoS attack scenario; to implement our proposal in the simulation software; and to explore the performance of our proposed approach in the simulated VSG scenarios.

This thesis proposes CognitiveCharge, a novel framework which enables independent nodes to make fully localised energy exchange decisions in real-time through careful weighing of the immediate and anticipated future energy and threat contexts. CognitiveCharge nodes make real-time decisions based upon immediate predictions of future conditions which increase their own utility and the utility of their self-identified spatiotemporal community. Individually, CognitiveCharge nodes continually monitor their own internal state as well as the behaviour of their neighbours through both direct interactions and passive observations. In addition to this, nodes participate in collaborative data exchange wherein they periodically disseminate aggregate information regarding themselves and their peers with immediate neighbours. Via a proposed suite of novel, predictive analytics, these metrics are combined to describe the perceived context for a given CognitiveCharge node for itself and its community. These analytics capture the complex interplay across multiple continuously changing dimensions, including network topology, energy resource behaviour, and trust and reputation dynamics. Combined together via real-time heuristic-driven decision-making, they allow for CognitiveCharge nodes to detect and react to unstable conditions, discovering and identifying suitably trusted opportunities for the localised exchange of energy.

1.5 Contributions

This thesis contributes to the field by providing a novel CognitiveCharge framework which enables VSG devices to perform fully-localised energy exchange in untrusted environments. The following key contributions together form CognitiveCharge:

- A suite of novel, cross-layer, predictive energy and threat context-aware analytics for capturing and interpreting the VSG environment from the perspective of a VSG node.

- A novel framework for analytics exchange, peer data integration, and real-time energy-resource utility-driven decision-making in untrusted VSGs.
- A proactive, collaborative peer testing mechanism for energy-resource behaviour evaluation in untrusted VSGs for detection and mitigation of EDoS attack.

The following publications have been the result of this work:

- M. Radenkovic and A. Walker, ‘CognitiveCharge: Disconnection Tolerant Adaptive Collaborative and Predictive Vehicular Charging’, in *Proceedings of the 4th ACM MobiHoc Workshop on Experiences with the Design and Implementation of Smart Objects*, in SMARTOBJECTS ’18. New York, NY, USA: ACM, 2018, p. 2:1-2:9.
- M. Radenkovic and A. D. Walker, ‘Contextual Dishonest Behaviour Detection for Cognitive Adaptive Charging in Dynamic Smart Micro-Grids’, in *15th IEEE/IFIP Wireless On-demand Network systems and Services Conference, WONS 2019, Wengen, Switzerland, January 22-24, 2019*, T. Braun, L. Lilien, and Z. Zhao, Eds., IFIP, 2019, pp. 44–51.
- M. Radenkovic, A. Walker, and L. Bai, ‘Towards Better Understanding the Challenges of Reliable and Trust-Aware Critical Communications in the Aftermath of Disaster’, in *2018 14th International Wireless Communications Mobile Computing Conference (IWCMC)*, Jun. 2018, pp. 648–653.

1.6 Thesis Structure

The remainder of this thesis is structured as follows:

Chapter 2 provides a comprehensive review of the literature related to this thesis' research focus. Chapter 2 begins by identifying and detailing a set of criteria important for approaches which seek to address combined self-organised energy resource awareness, P2P energy exchange, and adaptive resource-related security for vehicular smart grid environments. This is followed by a survey of existing works, predominantly focusing on techniques for aspects of adaptive energy resource awareness, energy threat context awareness, and reputation-based security in opportunistic networks.

Chapter 3 details how this is accomplished by detecting and reacting in real-time to dynamic in-network conditions across multiple layers, including mobility, resources, services, and security. An architectural overview of our CognitiveCharge solution is initially provided, which highlights the key challenges in the complex design space and the conflicting trade-offs fundamental to the problem domain. This is followed by detailed analytical and functional models of our CognitiveCharge framework, which provide a thorough explanation of the framework's components, including the algorithms underpinning our collaborative, adaptive, real-time, decision-making processes.

Chapter 4 presents the experimental methodology employed in order to perform rigorous analysis and evaluation of our CognitiveCharge proposal. Chapter 4 begins with a description of the simulation environment and an explanation of our use of hybrid real-world and pseudorealistic data traces to represent each dimension of the problem scenario. Each vehicular smart grid scenario is described, and a comparative analysis of the scenarios is provided. Following this, we detail the modelling of the energy depletion attack in these scenarios.

Chapter 5 concerns the rigorous evaluation of the implemented prototype of our CognitiveCharge framework utilising the experimental methodology detailed in Chapter 4. Chapter 5 begins by introducing the set of established and novel proposed performance measures used to thoroughly evaluate our CognitiveCharge proposal. Extensive evaluation of CognitiveCharge is then conducted in vehicular smart grid scenarios in which nodes are fully trusted, nodes are mutually untrusted, and when nodes are under active energy depletion attack by injected malicious actors. Under each of these conditions, measured performance characteristics for CognitiveCharge are compared against baseline conditions as well as benchmark and state-of-the-art works.

Chapter 6 discusses in-depth the broader context of the proposal presented in this thesis, taking into consideration the criteria set outlined in

Chapter 2. This chapter considers the feasibility and challenges of real-world deployments of CognitiveCharge in domains with varying constraints and dynamics. This includes deployments across alternative topologies with degrees of centralisation and heterogeneity. Furthermore, Chapter 6 provides a discussion of additional applications for aspects of our proposed CognitiveCharge framework.

Chapter 7 concludes the work presented in this thesis, summarising our CognitiveCharge proposal and highlighting the key findings of our analysis and evaluation in light of the research aims. The main contributions of the thesis are also outlined in Chapter 7. Finally, proposals for direct and indirect avenues of future research relating to this thesis are discussed.

Chapter 2
Related Work

2.1 Introduction

This chapter provides a comprehensive survey of state-of-the-art techniques for aspects of adaptive, self-organised security and energy resource exchange in opportunistic networking (OppNet) and vehicular smart grid (VSG) environments. The structure of the rest of this chapter is as follows. We begin by defining and motivating the set of criteria which establishes the lens through which the relevant existing literature is reviewed. These are derived directly from the research questions posed in Chapter 1. This is followed by a detailed description and discussion of relevant state-of-the-art approaches, categorised by each criterion. We then look at approaches utilising fully localised social analytics to inform real-time decision-making. Works relating primarily to relevant aspects of energy follow this. Localised energy resource awareness techniques and works exploring peer-to-peer energy exchange are subsequently explored. Next, literature investigating security for challenged networks and direct behavioural analysis mechanisms for the detection of malicious behaviour and derivation of first-hand trust are detailed. This includes indirect reputation management techniques for second-hand trust.

2.2 Criteria

A novel set of criteria is proposed which, together, comprehensively capture the fundamental facets of the problem scenario detailed in Chapter 1. The selected criteria scope the research focus and are applied in this chapter in reviewing state-of-the-art techniques for aspects of collaborative, adaptive, energy resource and security-aware approaches in heterogeneous vehicular smart grid environments. In addition to providing the lens through which existing work is reviewed, the following dimensions collectively delineate the solution space constraints and evaluation scope. A summary of each listed criterion is provided, with thorough detailing of the criteria set provided in the following subsections.

1. Fully Localised: Functionally effective in fully distributed and decentralised environments, independently making real-time decisions informed by partial knowledge.
2. Delay and Disconnection Tolerant: Capable of operating in challenged network environments with no assumption of end-to-end connectivity and in the presence of communication disruptions.
3. Social Encounter Awareness: Awareness of the topological dynamics of the constantly changing local network environment.
4. Energy Resource Awareness: Awareness of the state of local energy resources as well as the energy conditions of the broader network environment.
5. Adaptive Peer-to-Peer Energy Exchange: Ability to make decisions independently and adaptively regarding the direct attaining and offloading of energy with peers in the dynamic local network environment.
6. Contextual Threat Awareness: Ability to directly detect the misbehaviour of peers and assess the wider threat environment so as to adapt behaviour according to the risk of attack.

7. Adaptive Trust and Reputation Aware Decision Making: Autonomous decision making informed by the combination of both learned first-hand knowledge and collaboratively acquired second-hand peer information.

Opportunistic networks are intrinsically heterogeneous, distributed, and decentralised. As such, they are devoid of centralised, persistent control for local coordination amongst nodes. Fully localised contextual awareness and real-time decision making are therefore necessary prerequisites for anticipating and adapting to the manifold dynamics inherent to these complex temporal networks. Spontaneous, self-organising communities formed of independently collaborating nodes are crucial to facilitate provision of services atop opportunistic networks, e.g. in order to provide a viable energy exchange marketplace. In such environments, the degree of participation of each node is determined independently according to its own motivations. The variable internal objective functions driving behaviour are perpetually measured against the fluctuating advantages and disadvantages of participation. Collaborative networks thus occur when the interests of nodes are in alignment. For example, viable energy exchange networks are formed when the desires of nodes seeking energy complement those with energy looking to offload. Several core components underpin the ability to achieve fully localised self-organisation in highly multifaceted dynamic network environments:

- Fully Localised Contextual Awareness: Without centralised information dissemination, nodes must utilise localised environment sensing to inform future decision making. This can include measurements from direct observations of itself and others – such as from interactions with neighbours – in addition to any explicit data garnered via peer information exchange.
- Fully Localised Coordination: Coordination for distributed and decentralised systems inherently relies upon local, collaborative self-organisation. Varieties of clustering techniques facilitate fully localised coordination of cluster members.
- Fully Localised Decision Making: Real-time decision making driven by local contextual awareness and localised coordination. Without

centralised means, each node must be capable acting independently and autonomously making decisions which in real-time allow for nodes to adapt to the complex temporal dynamics of the network.

The mobility of vehicles in the vehicular smart grid paradigm naturally causes to disruptions network communications. DTN approaches are therefore fundamental to facilitating data communications and coordination of P2P energy exchange in VSG environments. Similarly, the mobile social nodes comprising OppNets, such as those intrinsic to VGSs, are tightly interwoven into human activity. Vehicles and IoT devices such as smartphones present clear examples of this. Not only do they increasingly carry and transmit sensitive information such as sensor-derived health data and usage but due to their nature they inherently follow human mobility patterns. The necessary criteria exposed here are twofold. Firstly, cross-layer social-awareness permits adaptive decision making which demonstrably improves performance in OppNets. Secondly, the privacy implications of large-scale service-oriented mobile social networks are significant and thus privacy-awareness is of concern. The socio-temporal graph concerns a given nodes regular and irregular contacts which are derived through multiple complementary encounter-based analytics such as contact duration and contact frequency. Together this allows us to describe complex spatiotemporal relationship dynamics. Combining various metrics to analyse node ties allows for identification and ranking of nodes in social clusters as well as those more transient 'vagabonds' [22].

As the behaviour of nodes in IoT networks is closely tied to human behaviour and these networks are inherently social, nodes should be fully aware of the distinct multi-natured, dynamic, ephemeral social graphs. For clarity we here highlight the difference between the social ego network and the social relationship network – both of which are important but represent distinct dimensions of the problem space. The social ego network describes nodes which come into direct contact with one another and their connectivity with respect to measures such as centrality and betweenness. The social relationship network is the nodes self-identified social ties such as friends and family. Whilst there may be crossover between the two, they are fundamentally distinct. A favourable node in the ego network may not necessarily be favourable in the relationship network and vice versa. Vehicular social networks can be broadly described as being dynamic, temporal communities of vehicles which share behavioural commonality such that they have recurrent spatial encounters [23].

In the context of delay-tolerant networking, mobility has been described as a ‘double-edged sword’ [24]. On the one hand, the mobility of nodes amplifies the volatility of network routes and, consequently, the likelihood of network disruptions; on the other hand, the physical movement of devices can benefit isolated nodes and communities by increasing the potential for access to contacts with sought-after resources. A range of techniques for network analysis which can be conducted both centrally and by independent nodes without access to a global network overview have been explored [25].

Energy resource awareness concerns nodes having an understanding of their independent socio-spatio-temporal energy state as well as an awareness of the energy state of their wider ego-network and the broader network as a whole. Energy is a fluctuating resource which nodes need to be aware of. P2P energy exchange is core to this work. The exchange of energy fundamental criteria and so approaches which facilitate this are important. Nodes must be able to take advantage of exchange opportunities based on their own awareness of the local and wider energy conditions.

Threat detection is fundamental to detecting and mitigating malice and malfunction. To exchange energy opportunistically in the dynamic VSG, nodes must be able to make real-time decisions based on the trust and reputation of their peers. This includes independent detection of malicious behaviour as well as trust and reputation exchange amongst peers.

2.3 Social Awareness

Large-scale networking environments with complex dynamic topologies benefit from exploiting social encounter-based analytics to inform aspects of decision-making. Localised, independent, real-time decision making calculated independently from the perspective of each node in order. Networks such as the VSG have predictable social dynamics and are non-random, exhibiting power law intermittent connectivity [24], [26]. Works reviewed in this section focus on social clustering approaches and social-aware protocols.

In opportunistic networks which exhibit social mobility properties, the efficiency of message routing can be improved by factoring social measures into decision making [26]. This is driven by independent analysis of encounter history and facilitated through collaborative, local peer information exchange via ‘summary vectors’ [27]. The ‘ego network’ concept describes a node's perspective of its local, socio-spatiotemporal neighbourhood [28]. A node’s ego network represents the complete set of past and present single-hop encounters – also termed the first-order zone [28]. Routing-related decision-making based on analytics over a node's ego network has been shown to achieve optimal and near-optimal performance in socially mobile environments [29].

Routing of data in opportunistic networks is largely dissimilar to routing in ad hoc networks due to the assumption of end-to-end connectivity in the latter case. Despite this, there are parallels between hop-by-hop forwarding mechanisms in opportunistic networks and on-demand route discovery mechanisms for ad hoc networking routing protocols. FRESH [30] is an efficient, omnidirectional blind search approach that improves route discovery by monitoring single-hop encounter history. Route discovery is accomplished via passive social analytics, specifically using the encountered neighbour with the lowest intercontact time for a given destination [30].

A number of routing protocols based upon ego network analysis have been proposed to enhance various performance characteristics of multi-hop routing protocols. These works primarily focus on improving aspects of message delivery, such increasing delivery success ratios, reducing delays, and minimising protocol overheads. PRoPHET [31], [32] is a probabilistic forwarding protocol that establishes adaptive, transitive delivery predictability based upon calculating the frequency of encounters between a PRoPHET node and each peer in its ego network. The MaxProp [33] protocol similarly utilises encounter frequency analytics but applies it in ranking the priority of undelivered messages for forwarding and dropping by estimating the likelihood of delivery. In delegation forwarding algorithms [29], [34] nodes are assigned utility measures which represent their routing quality for a message. Forwarding decisions are modelled as a form of optimal stopping problem [35], with routing costs reduced by forwarding messages to nodes of strictly increasing quality – calculated through implicit social metrics [34]. By combining semantic social relationship information (i.e. explicit labels [36]) with social encounter centrality measures to identify the popularity of nodes amongst detected communities, the BUBBLE [36], [37] protocol identifies how increased context awareness can achieve significant improvements in forwarding efficiency.

SimBetTS (2009) [38] demonstrates the efficacy of combining several lightweight utility measures of social networks for improving forwarding decisions in networks with attributes of social mobility. Similarity, betweenness, and tie-strength metrics are combined into a single utility which is compared against encountered neighbours in order to select for improving values. PRoPHET+ (2010) [39] extends the social-aware PRoPHET decision-making process to improve performance in constrained networks. This is accomplished by additionally considering dynamic weighted parameters for available bandwidth, buffer, power consumption, and popularity of nodes in the ego network. Many works build upon these core social clustering principles.

In recent works, social network analysis techniques have been extended from the topological plane and combined with characteristics of other dimensions. This allows for better describe a node's contextuality of decision-making. Recent research has explored utilising combined opportunistic vehicular and social communications for data processing, information processing and services and has shown that vehicles can collaborate over multiple dimensions and adapt to temporal dynamic networks. Café and CafRep [40] propose multi-layer adaptive congestion-aware protocols which combine social metrics with predictive analytics to direct network traffic away from congested areas of the network. Both protocols successfully reduce congestion whilst remaining considerate of resources and avoiding node overloading. CafRepCache [41] proposes a real-time heuristics and cross-layer analytics-based adaptive caching and forwarding approach. CafRepCache builds upon Café and CafRep, adding support for latency-aware collaborative caching. Café, CafRep and CafRepCache are all evaluated over diverse, dynamic temporal network topologies which include multiple real-world vehicular traces. E3F [42] proposes a cross-layer energy and network congestion-aware framework for resource-constrained opportunistic emergency networks, which avoids forwarding data to network regions which are not low on energy or likely to experience congestion. Using multi-dimensional real-time heuristics, E3F prioritises important nodes for protection against energy depletion and congestion while achieving high success ratios and improving node lifetime.

Real-world utilisation of some approaches are insufficiently adaptive due to the requirement for advance preselection of initialisation constants in a given environment, e.g. [31], [32], [39], [43]. This necessitates advanced analysis of the network environment to determine suitable values and it follows that the resultant performance is dependent upon the accuracy of values chosen. Reliance upon fixed values in decision-making limits the ability of the protocol to adapt to network changes, and so variable, unstable network environments can experience reduced protocol performance. This is overcome in works which are fully adaptive to the dynamics of the network and can consequently offer increased reliability in the presence of unstable conditions.

Surveys of social-aware routing, especially in DTNs (e.g. [44]) have highlighted a number of core open areas that are being actively addressed, including the metrics used to drive decision making (e.g. social ties), selfishness of nodes, adaptability and intelligence of protocols, and security and privacy. Whilst we do not address all of these, we build upon each area and apply these principles to a novel VSG scenario.

2.4 Energy Awareness

Peer-to-peer energy exchange describes the direct transference of electrical energy between two independent hosts at the edge of the smart grid, realising the ‘plug and play’ paradigm of the emerging ‘internet of energy’. This section presents a focused review of state-of-the-art approaches for distributed peer-to-peer energy exchange amongst heterogeneous mobile devices and static grid infrastructure. This literature review focuses on relevant and compatible approaches which satisfy our defined criteria.

As a result of varying system assumptions and research foci, several architectures for peer-to-peer energy exchange have emerged in the literature, primarily differing depending upon three key factors:

1. Assumed constraints for distribution of control, i.e. centralised versus decentralised coordination and decision making.
2. The degree of localisation of energy exchange, i.e. energy exchange facilitated only at dedicated sites (such as charging stations and car parks) versus energy exchange at any location.
3. The level heterogeneity of participating nodes, e.g. bidirectional vehicle-to-vehicle, unidirectional grid-to-vehicle.

The degree of distribution of energy exchange and the decentralisation of coordination represents two tangential dimensions across proposed systems for peer-to-peer energy exchange. Works can be classified by both the centrality of the coordination of the energy transfer and by their dependence on the electrical grid infrastructure for the physical exchange of energy. We include a range of works along these axes.

Amongst the most stringent of energy exchange approaches are those which are dependent upon both centralised coordination whilst simultaneously restricting energy exchange occurs at specific locations ('swapping stations'). The Danish EDISON project [45] was a joint academic and industry initiative to investigate vehicle-to-grid integration of a large electric vehicle fleet with the electrical grid for the benefit of both end users and the grid itself. The distributed EDISON electric vehicle virtual powerplant proposal allows the grid to accommodate the increased load of electric vehicles whilst taking advantage of local renewable energy resources to reduce CO2 emissions [45]. Energy and data flows in the EDISON electric vehicle virtual powerplant are vehicle-to-grid with charging of vehicle batteries being reliant upon semi-centralised scheduling allowing for only soft real-time local control [45].

Several works have explored optimising the scheduling of charging of electric vehicles from the grid infrastructure, e.g. [46] [47] [48] [49]. These approaches seek to account for both the needs and desires of independent electric vehicles whilst also optimising for centrally determined criteria. Such scheduling approaches assume existence of a trusted, centralised, controlling authority for coordination combined with a high degree of persistent connectivity for charge scheduling. This is not something that is present in our defined VSG scenario.

Mobile vehicle-to-vehicle energy exchange platooning using high energy capacity nodes such as buses and lorries is proposed in [50]. The approach in [50] uses cellular and ad hoc networking approaches to allow energy seeking nodes and needs only limited advance coordination in terms of the time the energy seeker expects to be connected. Despite this, the peer selection process in [50] requires complete manual user determination as to the suitability of any available charge opportunities and effectiveness depends strongly on accuracy of the route prediction mechanism.

A semi-distributed approach using energy swapping stations to facilitate vehicle-to-vehicle charging is proposed by [51] [52]. A vehicle-to-vehicle charging framework is proposed in [53] which uses a centralised location-based social network system to permit discovery of buyers and sellers so that electric vehicle owners can meet and conduct a financial exchange of energy. The social network system proposed by [53] requires users are fully active participants in the vehicle-to-vehicle energy trading process and have regular connectivity with the service in order be aware of seeker and supplier interests. Similarly, [51] proposes a single-period oligopoly game-based price control strategy which encourages electric vehicles with surplus to exchange energy to those with deficit by maximising discharging revenue and minimising charging cost.

Existing state-of-the-art approaches which only support a subset of peer-to-peer exchange (i.e. unidirectional grid-to-vehicle or vehicle-to-grid techniques) are insufficiently suited to fully opportunistic vehicle-to-vehicle energy exchange. Some avenues of research into peer-to-peer energy exchange are overly prescriptive and incompatible in addressing the core research question, e.g. those which require manual placement of CSs. Centrally optimising approaches are typically reliant on advance forecasting of energy pricing and demand for the coordination of charging. For instance, day-ahead energy pricing from and demand forecasting are widely used parameters in the formulation of daily charging schedules for electric vehicles, e.g. [54], [55]. Reliance on coarse-grained, advance knowledge to derive fixed, long-term plans results in inflexible solutions which cannot adequately respond to real-time behaviour which deviates outside the predetermined expected range. For example, at electric vehicle battery swapping stations, the number of batteries available for exchange is limited. Therefore, if a station overestimates the day's energy demands it can lose profit, having overspent on energy acquisition from the grid. Conversely, underestimating the total energy demand can lead to customers receiving batteries with insufficient capacity to carry out tasks – under some models leading to the receipt of compensatory discounts, e.g. [54]. To avoid the impact of prediction errors, the inclusion of error margins in model formulation are used by centrally optimising approaches to reduce the impact of inaccurate predictions. However, such techniques lead to suboptimal solutions when predictions are correct and can even exacerbate problems when the actual demand is unexpected. One strategy is for suppliers to acquire extra energy when charging to account for unpredicted increases in usage [55]. Whilst accommodating some flexibility in prediction accuracy, provision of extra energy is to the direct and indirect detriment of both energy suppliers and seekers. Over acquisition of energy results in an increase in associated costs in terms of charge time and overall quantity of energy required. These costs have a feedback effect which reduce overall access to energy as increased charge times lower throughput and higher prices are subsequently demanded for surplus to recoup losses. When demand is underestimated the swapping station will have an energy deficit and when overestimated the swapping station loses profit.

In terms of energy awareness and integration of energy-related behaviour into real-time decision making, a large number of recent works (e.g., [56], [57], [58], [59], [60], [61]) seek to minimise energy consumption of the underlying DTN network protocols and associated computation cost. Whilst there is room for future works to investigate this in greater detail, we consider the energy consumption of the underlying communications protocol and any related computation to be out of scope of this thesis. This thesis focuses on actively roaming EVs where the energy consumption through physical movement is substantially higher than the energy consumption through computation associated with data communications and algorithmic decision making. EVs in the VSG scenarios central to this work can acquire and offload substantial amounts of energy. The majority of works looking at energy consumption in this area are focused on low-energy wireless devices which might rely on low-voltage, intermittent power supply (e.g. solar) and have limited battery storage capacity. The aim is therefore typically to sustain some level of service provision whilst minimising energy expenditure. Whilst these works are certainly energy aware, the underlying techniques do not directly apply to this work in terms of energy exchange and energy resource saving as we consider energy expenditure by nodes to be unavoidable and fundamental to a node performing it's primary function. For example, we consider an EVs core use to be in driving and transporting people and goods. Unlike works which look at aspects of optimisation and scheduling of charging (such as [62], [63], [64], [65]), we consider that the vehicles must be able to expend energy at any and all times, provided that there is sufficient battery to do so. We embrace principles from works which monitor and process consumption of energy but, in this work, we do not consider integration of decision-making processes which can result in a lack of full, complete and always immediate access to all available energy. Nevertheless, there is certainly scope in future work to bring these areas together to further refine and improve energy usage.

Automated energy transfer from a dedicated supplier node to energy seeking nodes has been considered by works such as [66], [67], which have explored usage of roaming auxiliary devices to supply energy to stationary nodes in need of charging. Although on a smaller scale and focusing on 'topping-up' batteries attached to wireless sensor nodes, the idea presents as an automated version of the EV emergency chargers currently offered by roadside assistance and automotive services companies such as the RAC [68]. Both the research and the commercial offerings in this respect focus on a strict hierarchical supplier-seeker relationship. Fundamental to this work is the flattened, dynamic relationship where nodes are universally considered 'prosumers' and can act as energy supplier or seeker depending on their own context and can change roles or even present as different roles to different peers.

2.5 Threat Awareness

Collaborative techniques in opportunistic networks are vulnerable to a range of attacks (including denial of service) if data received from peers is blindly trusted [69]. An adaptive and flexible k-anonymity approach for opportunistic networks is proposed in [70], which uses collaborative cross-layer heuristics to allow for nodes to identify and choose suitable anonymisation overlay nodes for use without degradation of network performance. Extending [70], OCOT-AA [71] adds a peer testing mechanism which allows for pre-emptive analysis of obfuscation behaviour. OCOT-AA [71] efficiently and rapidly detects misbehaving nodes when a trusted peer is available for collaborative peer behavioural analysis. Misbehaviour is identified via a collaborative peer testing scheme. CogPriv [72] builds upon [71] to provide a real-time, fully-localised privacy framework for personal clouds which uses collaborative smart probing and multi-service isolation via a virtualisation layer for adaptive, on-demand privacy in heterogeneous networks.

Black hole attacks are a similar form of DoS attack in mobile networks such as mobile ad hoc networks (MANETs) and delay-tolerant networks (DTNs). Many approaches to black hole detection and reputation management have been explored in the context of wireless networks using a wide range of approaches. For instance, machine learning [73], fuzzy logic detection [74], game theory [74], the TEAR mechanism for hotspot avoidance [75]. Many works extend the existing protocols to capture information about node behaviour and to identify mis-behaviour, e.g. [76]. Whilst there is a wealth of research exploring approaches to mitigation to black hole attacks in these environments [77], [78], [79], the fundamental differences between energy and data limit the efficacy of many approaches as being applied to EDoS attacks in VSGs.

Blockchain strategies provide a mechanism for threat avoidance through dissemination of certified records. By checking these records, it is possible to determine whether a node is falsifying local information. ‘Off-chain’ transactions present a promising approach wherein a transaction between two nodes does not need to be logged to the blockchain or can be deferred until later, reducing reliance on centralised and time-bound communications. In [80], authors explore centralised, decentralised, and hybrid architectures for smart contracts based upon blockchain systems. Many approaches to off-chain transactions rely on trusted third parties, e.g. for escrow. An alternative approach makes use of secure payment channels and cryptographic algorithms (such as the Hashed Time Lock Contracts employed in the Lightning Network [81]) to allow for revocable P2P off-chain payments. Despite providing secure off-chain transactions, an initial commitment transaction must still be broadcast to the blockchain before off-chain trading can begin [81]. Whilst cryptographic approaches, where possible, offer strong security guarantees, they do not account for malicious behaviour of authorised nodes which later present a threat (e.g. a negative behaviour change due to being hacked). There is an assumption in typical blockchain technologies that the actor authorising participation in the exchange is the owner or authorised operator of the node, which may not be the case if the device under the influence of a malicious entity. We therefore make use of trust schemes and reputation techniques which are known effective countermeasures in such environments.

Emerging distributed ledger technologies such as the blockchain have been explored by recent works for enhancing both security and end-user privacy in untrusted, distributed, and decentralised internet-of-things systems for a diverse range of environments, including smart energy, and healthcare, amongst many others. Viability of the blockchain has been considered by many recent industry and research projects for financial energy trading in smart microgrids. Blockchains are expected to see wide deployment in the internet-of-things due to their inherent support for transaction immutability and provision of cryptographic assurances in decentralised environments [82]. Blockchain deployments are categorised as either public, private, or consortium with various algorithms employed for distributed consensus, the most common being variants of proof-of-work, proof-of-stake, or proof-of-activity. Privacy of blockchain transactions uses emerging techniques such as address anonymisation, transaction privacy, and smart contracts [82].

Despite the advantages of blockchains there remain open challenges in meeting the performance, scalability, and data consistency requirements needed for future deployments [80]. Eventual consistency in distributed systems guarantees that once updates to a data entry have ceased, all retrievals of the recorded entry will eventually yield an identical response. Providing stringent consistency guarantees is beyond the scope of our work. We assume existence of a distributed ledger or blockchain deployment for recording financial transactions of energy. Our approach compliments existing blockchain approaches. The specifics of the deployment (e.g. architecture, consensus algorithms, etc.) do not impact our proposal architecture or effectiveness as we assume the most stringent and restrictive case for each (e.g. our techniques do not require that transactions on the blockchain are readable and therefore emerging privacy techniques can be seamlessly integrated with our approach.

The Brooklyn Microgrid is an ongoing, private blockchain-backed energy marketplace comprising residences and local solar renewable energy resources ran by the LO3 Energy company in New York, USA [83]. Buying and selling of energy is automated based upon user specified criteria and in emergency situations the Brooklyn Microgrid can operate whilst fully disconnected from the wider grid [83]. Though demonstrating the viability of blockchain energy marketplaces in microgrids, the Brooklyn Microgrid is completely static and energy flows are coarse-grained; all energy trading is time synchronised and not real-time [83].

Specifically targeting renewable energy, the NRG-X-Change [84] protocol proposes a scalable, blockchain-backed marketplace which incentivises energy exchange between local prosumers and the grid. Locally generated energy supplied to the grid earns prosumers NRGCoin [85], a virtual cryptocurrency analogous to Bitcoin. Smart metering agents over time determine strategies for buying and selling of energy as well as for non-immediate conversion of NRGCoin to more widely accepted currencies such as fiat money (e.g. pounds sterling, US dollars) [84] [85]. Feasibility of the both NRGCoin and the NRG-X-Change platform was demonstrated in a static environment in [86]. Despite insufficiently real-time. Responsiveness of such an approach with increasingly dynamic scenarios and unpredictable events.

PETra [87] is a privacy focused approach to energy transactions for prosumers within smart microgrids which share a common energy link with the wider grid. Participating microgrid nodes anonymously negotiate energy trades in advance via a semi-distributed bidding service with each microgrid having a controller to predict usage and dictate pricing [87]. PETra's security architecture assumes that smart meters can generally be trusted and therefore requires tamper resistant hardware be deployed to prevent theft or malicious behaviour [87]. Similarly, PETra is reliant upon trusted, semi-centralised operators and requires trade decisions be made in advance without necessitating proof of real-time local energy availability [87].

PriWatt [88] proposes a reliable, privacy oriented proof-of-work-based blockchain-backed energy trading platform which permits anonymous price negotiation and replication of transactions to avert failure and attack. The PriWatt architecture is a hybrid of hierarchical, centralisation and peer-to-peer networks [88]. Although PriWatt is not reliant on central pricing controls, it requires availability of distributed system operators for mediating transactions and additional communication with an 'auction board' for advertising of anonymous supply and demand requests [88]. Security of PriWatt energy transactions depends upon the ability of both buyer and seller to communicate in real-time with multiple common entities [88].

A smart contract-based permissioned blockchain for secure charging and reputation based delegated Byzantine fault tolerance consensus algorithm are proposed by [89] in which centralised aggregators manage local energy resources and coordinate the charging of collocated electric vehicles from the grid. Three adversarial groups are presented in [89], namely malicious energy providers, malicious, energy consumers, and malicious trusted third parties. Use of permissioned blockchain means only pre-authorised nodes are trusted to process transactions, limiting flexibility.

PETCON [90] is a semi-centralised, privacy preserving P2P energy trading system which uses a consortium blockchain approach wherein local, preselected aggregator nodes act as energy brokers and peer coordinators, controlling the logging of transactions to the blockchain via proof-of-work consensus. Authors highlight the importance of incentive schemes for balancing the supply and demand amongst heterogeneous collocated electric vehicles and additionally propose an auction mechanism for energy price negotiation [90]. The PETCON proposal requires that nodes in advance commit to a certain amount of charge-time and establish a fixed price in advance of any possible transaction. Furthermore, nodes who wish to participate in energy exchange must use an available local aggregator to handle coordination. Vehicle-to-vehicle energy exchange in the PETCON system is therefore not fully opportunistic and is insufficiently adaptive to real-time dynamic conditions. Other protocols such as EBR [69] resort to a mechanisms for digitally signing data.

Security, in particular, is an open challenge in the area of DTNs and OppNets [91], [92]. In this work, as highlighted, we consider behavioural analysis as the core mechanism for threat detection suitable for the defined VSG scenario. In the scope of the VSG, as highlighted, we further consider reputation and peer trust dissemination as crucial for spreading awareness of malicious behaviour and mitigating the impact of attacks. There are range of strategies for threat-awareness information dissemination via trust and reputation management schemes, amongst others with a focus on data routing and forwarding (e.g. [75], [75], [93], [94], [95], [96], [97], [98], [99], [100], [101], [102]). Broadly, the architecture of these are similar in necessitating some initial trust value derivation, providing a strategy for reputation dissemination, and then adjusting the trust value using reputation ratings received from peers via a local trust update mechanism. In the context of OppNets, real-time decisions are made alongside routing related decisions as to whether a peer is trusted enough to participate in an exchange of data.

In depth consideration of trust and reputation management strategies is beyond the scope of this work. Whilst we apply fundamental principles to manage trust and reputation in this work, there is nevertheless room for future work to pick up on this in greater depth to further improve accuracy and resilience of the proposed work to attack. We focus on a single tier threat detection and reputation mechanism to meet the core aim of this work. However, for a practical implementation in the real-world it would be essential to broaden the scope of this work to consider a broader range of attacks, as well as attacks against the detection mechanism themselves (such as ballot stuffing) and the necessary layered defences to thwart these.

Chapter 3
CognitiveCharge

3.1 Introduction

This chapter details our novel CognitiveCharge proposal. CognitiveCharge enables predictive energy availability and threat context awareness to provide real-time identification and exploitation of suitably reliable opportunities for P2P energy exchanges. CognitiveCharge nodes monitor local and regional information from first-hand observations and second-hand collaborative information propagation to capture the dynamically changing socio-spatio-temporal energy and security contexts. Via a novel suite of multi-dimensional, complementary, real-time, predictive analytics, CognitiveCharge nodes can make adaptive, heuristic-driven decisions for on-demand energy acquisition and supply offloading in disconnection-prone, distributed, and decentralised VSGs. By adaptively balancing dynamic energy and contextual security sensitivity, CognitiveCharge nodes make real-time decisions in highly dynamic untrusted networks.

This section has briefly introduced our CognitiveCharge proposal. The remainder of this chapter is structured as follows. Firstly we detail the architecture of our CognitiveCharge proposal. Following this, we provide the formal multi-dimensional spatiotemporal graph definition which we use in this work to model both the global network and nodes' independent ego networks. A minimal network example is then given to clarify aspects of the model and highlight the perspective differences.

3.2 Architecture

Building on established works (including [40], [41], [103], [104]), CognitiveCharge combines data from multiple layers. These layers are monitored first-hand in addition to integrating second-hand propagated information. As in similar works in this area, we distinguish between core metrics – for instance, monitored data such as energy consumption through usage and social tie strength – and analytics – the predictive 'wrappers' which build upon the metrics and capture the state of the network to inform future decision making. Using energy level as an example, in this work the raw energy level over time would be considered a metric. Applying a prediction over the set of recent energy levels, such as to anticipate when a battery might deplete, would be considered an analytic. Analytics are therefore composed of metrics. In other works, analytics and metrics be also termed 'statistics' or 'properties'. To help illustrate this, Figure 4 shows the hierarchy of metrics and analytics in an example of a decision-making engine. This is essentially a scaled-back and simplified version of CognitiveCharge. In OppNets, this form of decision-making would drive data routing decisions when nodes encounter one another. Here, we look at a hypothetical Simple Offload Decision Engine (SODE) which might be used to make a decision about forwarding or receiving energy. We can imagine that each node in a hypothetical VSG scenario is equipped with SODE and wants to know, when it encounters another node, whether it has surplus energy that it can acquire or offload or whether it should hold on to its remaining energy for the time being. The battery level metric tracks the battery level over time and the depletion analytic wraps around this to monitor the rate at which the level is increasing or decreasing. Similarly, the inter-contact time metric measures the time between encounters, with the encounter analytic predicting the next time the node will have an opportunity to offload or acquire energy. The SODE component of each node therefore balances the immediate and forthcoming needs of the node, considering the anticipated opportunities for exchange. This is a contrived example, purely to illustrate and clarify key concepts for this work in a field where terminology is very overloaded. For CognitiveCharge, and similar works in the literature which model decision making in this way (e.g. [41]), analytics are composed of multiple metrics and seek to capture a facet of the network and multiple analytics inform the decision-making processes. This, for CognitiveCharge, means continually monitoring data across layers, both first and second hand, so that meaningful analytics, which provide higher-level context, can anticipate key events to be fed into the decision engine for real-time decision making.

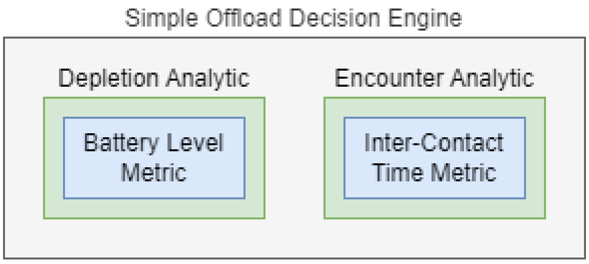

Figure 4 Example Decision Engine Highlighting Metrics and Analytics

The distinct layers which comprise the high-level architectural overview of our conceptual model are shown in Figure 5, which highlights: the physical, data networking, and peer-to-peer energy exchange layers; the social encounter dynamics layer; the energy resource need and availability layer; the trust and reputation-based security layer, and the supply and demand energy marketplace layer. Each of these layers is represented as a dynamic temporal graph which is continually evolving as a result of internal and external influences across multiple dimensions. By combining locally calculated predictive analytics from metrics monitored for each dimension with exchanged peer information, each independent CognitiveCharge node can make informed decisions based on the current state and predicted future state of each layer. CognitiveCharge adaptively manages trade-offs between these dimensions through multiple predictive real-time analytics.

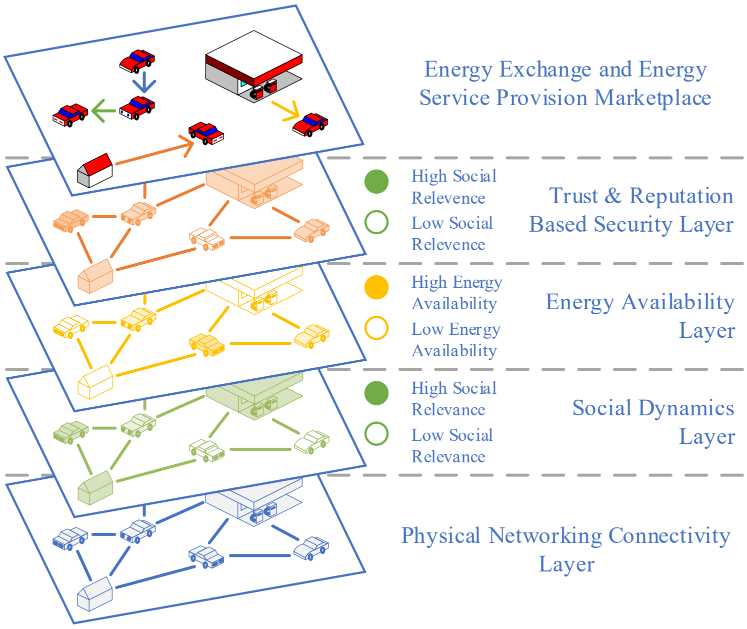
Energy Exchange and Energy Service Provision Marketplace
High Social Relevence
Low Social Relevence
Trust & Reputation Based Security Layer
High Energy Availability
Low Energy Availability
Energy Availability Layer
High Social Relevance
Low Social Relevance
Social Dynamics Layer
Physical Networking Connectivity Layer

Figure 5 Architectural VSG Layers

At the physical layer is the dynamic network topology wherein independent nodes spontaneously form local networks through temporal encounters resulting primarily from node mobility. Mobile nodes independently roam the environment, per their primary behaviour (e.g. as taxis, buses, or consumer vehicles), interact locally with one another as well as with static nodes (e.g. traditional charging infrastructure at dedicated refuelling stations and charge-points installed at homes). Dual connections can be made between collocated nodes at the physical networking layer, representing the information communication and energy flow channels, respectively. The temporal networks formed from these connections can correspond one-to-one if the range of each is identical, otherwise, the range of energy connectivity is assumed to be less than that for data. V2V and V2G connections in these environments are short-lived and frequently changing. There is no assumption of end-to-end path connectivity and exchanges of both information and energy are required to be both real-time and localised.

The social, temporal graph layer concerns a given node's regular and irregular contacts which are derived through multiple complementary encounter-based analytics such as the contact duration and contact frequency. Together, these social analytics allows us to describe complex spatiotemporal relationship dynamics between nodes and node groups. Combining various metrics to analyse node ties allows for the identification and ranking of nodes in social clusters as well as those more transient ‘vagabonds’ [22].

Atop these layers is the fluctuating geo-temporal availability of energy. In vehicular smart grids, energy is transient. Predominantly, energy moves fluidly as it is expended and exchanged amongst mobile vehicles and the grid. Availability can also be more sharply disturbed by unplanned events such as grid outages resulting from attack or malfunction. Even the existing electric vehicle infrastructure (i.e. vehicles only acquiring energy from static grid CSs) has geo-temporally dynamic energy availability as charge-points can be broken, undergoing servicing, or experiencing blackout.

The trust and reputation-aware security later captures the perceived risk of attack from peer nodes via detected and predicted malicious behaviour. Amongst socio-spatiotemporal network communities, trust and reputation may be incongruent as perspectives of collated reputation values differ between independent CognitiveCharge nodes. Node with strong mutual trust may not have the same opinion of a common peer, either because they have received incompatible second-hand opinions or have had different first-hand experiences with the peer; for example, due to malicious behaviour only targeting specific nodes. Therefore, for each node, its perception of its peers’ reputation changes spatiotemporally as attacker behaviour may be nonuniform, e.g., targeting groups of nodes or geographic regions. Finally, at the uppermost layer there is real-time supply and demand of energy and provision of energy consuming services, such as mobility and exchange.

3.3 Model Formulation

As detailed in previous sections, the VSG scenario central to this work is highly dynamic and multi-dimensional. In this section, we present our model formulation of the VSG. This model is used consistently throughout this work however it is essential to note that there are critical distinctions between the global system overview and the perspectives of independent nodes within the system. The former represents a complete but external birds-eye view of the multi-dimensional spatiotemporal graph. This global model is extremely useful for analysis. We first give the general, global graph model. Subsequently, we provide further model assumptions specific to the VSG scenario central to this thesis. An example VSG network scenario is provided later in this chapter to illustrate further how the global and ego network graph models presented here relate to one another.

To formally model the complex dynamics fundamental to our spatiotemporal cross-layer scenario, we build upon established temporal graph models as presented in [105], [106]. We extend these models to suitably capture the high dimensionality of the problem space and VSG context. Therefore, we model the network as a multi-dimensional, time-dependent digraph with multiply weighted edges G where V is the set of vertices, E is the set of edges, D is the set of all dimensions, C is the set system (family of sets) of valid weights (capacities), and finally, T is the complete set of possible times which appear as interval bounds in edge tuples. Our graph model is shown in Equation 1.

$$
\begin{aligned}
V &= \{v_1, v_2, \dots, v_m\} \\
E &= e_1, e_2, \dots, e_n \\
D &= \{d_1, d_2, \dots, d_o\} \\
W &= \{W_1, W_2, \dots, W_p\} \\
T &= \{t_1, t_2, \dots, t_q\} \\
G &= (V, E, D, W, T)
\end{aligned}
$$

Equation 1 Graph Model

Each edge of the graph is represented as a 7-tuple $e = (u, v, d, w, c, t_i, t_j)$, which indicates a directed connection from node u to node v in dimension d with weight w of possible capacity c at between time intervals t_i (the start time) and t_j (the end time) $[t_i, t_j]$. The duration of the connection can be obtained simply as $t_j - t_i$. As such, the set of possible edges E for a VSG scenario is a strict subset of the Cartesian product of the set of possible values for each tuple element $E \subseteq V \times V \times D \times \bigcup W \times \bigcup W \times T \times T$. This set may be further reduced to omit invalid edges, such as by ensuring that there are no conflicting or invalid weight and capacity combinations in a given dimension

($d_i \wedge w, c \in W_i \wedge w \leq c$ for each edge), that time intervals are possible ($t_i < t_j$), and there are no loops ($u \neq v$).

In this multi-dimensional graph model, each edge includes the capacity value representing the maximum possible link connection between any two nodes and the weight representing the actual usage. The capacity value is determined by the physical constraints of the network as well as any subsequent constraints that the nodes themselves may apply. As such, every temporal connection in the graph has a weight w which is bound by a link capacity $w \in C \subseteq W$.

To simplify notation for operations over graphs and their constituents, we adopt and adapt some common functional mechanisms to denote operations over sets conveniently. For example, to extract snapshots of the multi-dimensional time-dependent graph G by filtering the edges based on some desired criteria and accessing specific elements of tuples. As shown in Equation 2, the function $\pi_n(x)$ accesses the nth projection of a given tuple x. For instance, $\pi_2(a, b)$ would return b. To simplify the notation in the context of this model, we also utilise the tuple variables defined above when projecting named elements onto the 7-tuple edges, i.e. $\pi_d(e) \equiv \pi_3(e)$.

$$\begin{array}{rcl} \pi_i\colon X_1 \times \cdots \times X_i \times \cdots \times X_n & \longrightarrow & X_i \\ (x_1, \ldots, x_i, \ldots, x_n) & \longmapsto & x_i \end{array}$$

Equation 2 Projection Function

As shown in Equation 3, a generic, higher-order function for filtering an arbitrary set X into a subset $X' \subseteq X$ based on some predicate $p\colon X \to \mathbb{B}$ can be defined such that $\forall_{x \in Y}(p(x) \wedge x \in X)$ where ${X'}^X$ denotes the set of all functions from X to set X'. The power set is denoted as usual in the domains and codomains of these functions as $\mathcal{P}(X)$. To illustrate this, the edges E' of the one-dimensional temporal subgraph of graph G in dimension d_i that occur strictly within the time interval $[t_i, t_j]$ can then be obtained from the set of all edges E as $E' = \text{filter}\ (e \mapsto \pi_3(e) = d \wedge t_i \leq \pi_6(e) \leq \pi_7(e) \leq t_j)\ E$. We can similarly define the well known higher-order map and reduce functions. The map function (Equation 4) applies some operation to each element of a set X to return a new set Y with the same cardinality. These functions are composable such that $g(f(x)) \equiv (g \circ f)(x)$.

$$\begin{array}{lcc} \text{filter:}\, \mathbb{B}^X & \longrightarrow & \mathcal{P}(X)^{\mathcal{P}(X)} \\ \quad p & \longmapsto & (X \mapsto \{x \mid p(x) \wedge x \in X\}) \end{array}$$

Equation 3 Filter Function

$$\begin{array}{lcc} \text{map:}\, Y^X & \longrightarrow & \mathcal{P}(X)^{\mathcal{P}(X)} \\ \quad f & \longmapsto & (X \mapsto \{f(x) \mid x \in X\}) \end{array}$$

Equation 4 Map Function

For any single-dimensional graph snapshot, a square adjacency matrix representation $\boldsymbol{A}$ can be obtained from a set of edges E via the function $\mu(e)$ (for $e \ \in E$). For a matrix $\boldsymbol{A}$ (or equivalent set representation), the element at the ith row and jth column can be accessed via the function $\alpha_{i,j}: X^{m \times n} \to X$, where $X^{m \times n}$ denotes the set of all m by n matrices comprising elements of X, for which we use the shorthand $\boldsymbol{A_{i,j}}$. The 'index of' function $\iota: V \to N^+$ maps a node $u \ \in V$ to a unique natural number satisfying $\exists!\, u \in V, i \in [1..|V|](\iota(u) = i)$. In the adjacency matrix representation of a network, $A_{i,j}$ is set to weight w if there exists a direct connection from node u_i to node u_j, otherwise zero (or some other suitable value if zero would conflict with a valid graph weight, e.g. $\phi \notin W$. For a simple adjacency matrix where $\forall i \in [1..m], j \in [1..n]\big(\boldsymbol{A}_{i,j} \in 0{,}1\big)$, it is well known that the number of walks of length (hops) p between two nodes i and j can be derived by raising the adjacency matrix to the power p, i.e. $A^n_{i,j}$.

The global network model facilitates understanding of the cross-layer spatiotemporal events as they actually occur. Whilst useful for understanding the network, the VSG scenarios central to this work do not permit for global knowledge-based decision making as there is no central authority with persistent connectivity to all nodes. Partial knowledge combined with a mutual lack of trust necessitates localised decision making. The local ego network perspective considers the view of the network from the perspective of a single node. Unlike the global graph model, which presents the true state of the system, the ego network is localised and context-dependent for each node. As it is not possible for nodes in the highly dynamic and disconnection-prone VSG to have a global network overview, two separate nodes will most likely have different and possibly conflicting ego networks. This is particularly true in the untrusted VSG contexts of this work where mutual trust must also be obtained.

The ego network perspective [38] is extremely important because it represents the view of an individual node of its immediate neighbours. This is a crucial perspective to consider when considering how independent CognitiveCharge nodes analyse their environment and make real-time decisions based on partial knowledge derived from first and second-hand information. Complementing the global graph model, the ego network model allows us to understand the local dynamics and interactions between nodes before, during, and after network events.

The ego network perspective of the network represents the view of an individual node of its immediate neighbours. It is important to note that in this context we consider the neighbourhood of a node in a socio-spatio-temporal sense and so a node is a neighbour of another providing there has been a single-hop encounter between them at some point. We can find the neighbourhood of a node from the global set of edges for a given node u as $E' = \text{filter}(e \mapsto \pi_1(e) = u \vee \pi_2(e) = u)E$.

3.4 Network Example

To more clearly illustrate our multi-dimensional socio-spatio-temporal graph model, a minimal example network G is provided. Recall that we define each edge is a 7-tuple $e = (u, v, d, w, c, t_i, t_j)$. Figure 6 and Equation 5 show a representation of the time interval snapshots of G for the two dimensions d_0 and d_1. From the perspective of the overall network, the edges of G are given by E, the nodes $V = \{u_1, u_2, u_3\}$, dimensions $D = \{d_1, d_2\}$, valid weights $W = \{\mathbb{N},\ \mathbb{N}\}$, and time intervals $T = \{1, 2, 3\}$.

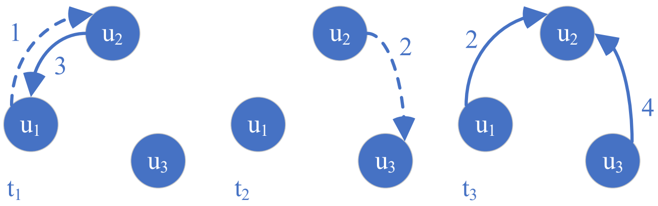

Figure 6 Network Example

$$E = \{(u_1, u_2, d_2, 1,3,1,2), (u_2, u_1, d_1, 3,3,1,2),$$
$$(u_2, u_3, d_2, 2,5,2,3), (u_3, u_2, d_1, 4,4,3,4),$$
$$(u_1, u_2, d_1, 2,5,3,4), \}$$
$$G = (V, E, D, W, T)$$

Equation 5 Network Example

Using the defined model, we are able to convey and extract a significant amount of multi-dimensional, spatiotemporal information about the network as represented by graph G. The set of edges $E' \subseteq E$ representing the subgraph of G at time interval [1,2] can be extracted using the method shown previously as $E' = \text{filter}\,(e \mapsto 1 \geq \pi_5(e) \leq 2 \wedge 1 \geq \pi_6(e) \leq 2)\,E$. We can calculate a total deficit in network utilisation across dimension d_2 by filtering where weight is less than capacity and summing the difference between the two for each edge.

$$\sum \text{map}\,\big(e \mapsto \pi_4(e) - \pi_3(e)\big)\ \text{filter}\,\big(e \mapsto \pi_4(e) < \pi_5(e)\big) \circ \big(e \mapsto \pi_3(d_2)\big)\ E$$

Furthermore, we can explore the ego network perspective of each node. To illustrate this, consider the case of obtaining a matrix containing information regarding the incoming usage of links in dimension d_1 from the perspective of node u_2.

$$\boldsymbol{A} = \mu \circ \text{filter}(\pi_3(e) = u_2 \wedge \pi_3(e) = d_1)E$$
$$\boldsymbol{A} = \begin{bmatrix} 0 & 2 & 0 \\ 0 & 0 & 0 \\ 0 & 4 & 0 \end{bmatrix}$$

3.5 Model Context

For our VSG scenarios, we consider a set of dimensions corresponding precisely to the architectural layers detailed previously. Thus, the complete set of dimensions can be given as $D = \{d_D, d_E, d_S, d_A, d_T, d_M\}$ for the highlighted real-time dynamic layers of physical data networking (d_D), physical energy exchange connectivity (d_E), social dynamics (d_S), energy resource availability (d_A), trust and security (d_T), and the energy trading marketplace (d_M).

All edge weights in this work are normalised to values in the closed unit interval $[0, 1]$ to help reduce the complexity of any cross-layer edge calculations. To explore our proposal with the strictest VSG configuration, in this work we consider physical data networking and energy exchange connectivity to be both bidirectional and unicast. It is important to note that this does not preclude multicast data or energy exchange from working with our proposed CognitiveCharge approach. Rather, this restriction represents a stringent core of the VSG definition in line with existing real-world mechanisms. By making this assumption with our CognitiveCharge model, we maximise the utility and relevance of our results to more flexible VSG scenarios.

At the cost of increased memory overheads and computational processing for analysis of graphs, to capture all cross-dimensional connectivity and node interactions in our VSG scenarios in practice we can assume a realisation of this model with consistent connectivity sampling at a suitably low fidelity that is below the duration of the smallest connection. This assumption is necessary for the highly dynamic real-time network model fundamental to this work, as we are particularly interested in localised events which coarse-grained connectivity snapshots could miss. Consequently, the overall representation of graph G is expected to be highly disconnected and formed of many socio-spatio-temporal subgraphs.

3.6 Self-Organised Coordination

Our proposed clustering protocol for CognitiveCharge nodes provides lightweight coordination mechanisms for contextual information dissemination and local energy exchange. These mechanisms are intended to operate atop, and be integrated with, any functionally compatible set of protocols. We therefore do not specify a precise format for messages, nor do we stipulate the underlying communication protocols. This section further highlights the scope of our cluster-based coordination mechanisms as well as presenting some key areas for consideration.

We presume that communication amongst CognitiveCharge nodes is private and our approach does not require that nodes have the ability to eavesdrop on messages that are not directly addressed to them, even if acting as an intermediary hop. In order to be maximally compatible with existing privacy-focused approaches, we presume that CognitiveCharge nodes can only ascertain information which is explicitly addressed to them or passively observed. Similarly, CognitiveCharge nodes do not rely on overheard communications metadata in order to be compatible with techniques which use peer forwarding for privacy and anonymity, e.g. [70], [71]. Specification of mechanisms to achieve message content privacy is beyond the scope of this thesis however there are existing approaches which could be used to support this. In particular, cryptographic techniques for distributed and decentralised networks.

In energy interest availability clusters, the cluster head continually adapts the price at which it is prepared to sell energy based upon a wide range of criteria including both its willingness to sell and the current local demand. The price of energy is therefore implicitly informed, in part, by peer behaviour such as the fluctuating demand for energy. An extension to this mechanism could consider explicit pricing feedback amongst buyers and sellers and our outlined availability clustering approach could be integrated with mechanisms for energy pricing negotiation mechanisms. Such strategies are beyond the scope of this work however several proposals exist which can be incorporated into our CognitiveCharge proposal, for example, multi-party bidding and auctioning approaches.

Figure 7 illustrates our three-tiered coordination clustering mechanism using an example network comprised of five nodes, as shown from the perspective of the energy supplying node n_2. The three, overlapping clusters of real-time energy exchange, energy availability interest, and ego network community highlighted in Figure 3 are represented by the sets $X = \{n_3\}$, $Y = \{n_1, n_3\}$, and $Z = \{n_1, n_3, n_4\}$, respectively, with the relationship $X \subseteq Y \subseteq Z$ holding in all cases. Figure 7 further shows how the levels of our hierarchical cluster-based coordination mechanisms correspond to the time and location criticality of the decision-making being made by CognitiveCharge nodes, with V2X energy exchange being the most highly time sensitive.

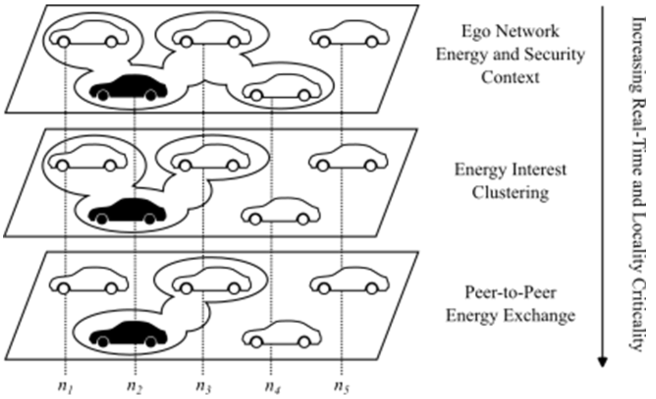

Figure 7 Tiered Self-Organised Clustering

We propose a lightweight, interest-based, socio-spatio-temporal clustering mechanism through which CognitiveCharge nodes can collaboratively coordinate real-time, opportunistic exchange of energy. Energy interest-based clustering connects a single energy supplier (a node with a surplus of energy that it is willing to offload) with nodes in its dynamic, immediate, single-hop neighbourhood which wish to acquire energy. At the head of each energy interest-based cluster is the energy supplier which initiated clustering. The presented mechanism is entirely proactive – nodes must opt-in as a cluster member in order participate. An illustration of this is provided in Figure 8 which shows a local, temporal single-hop network, in particular, the ego network of the supplying node.

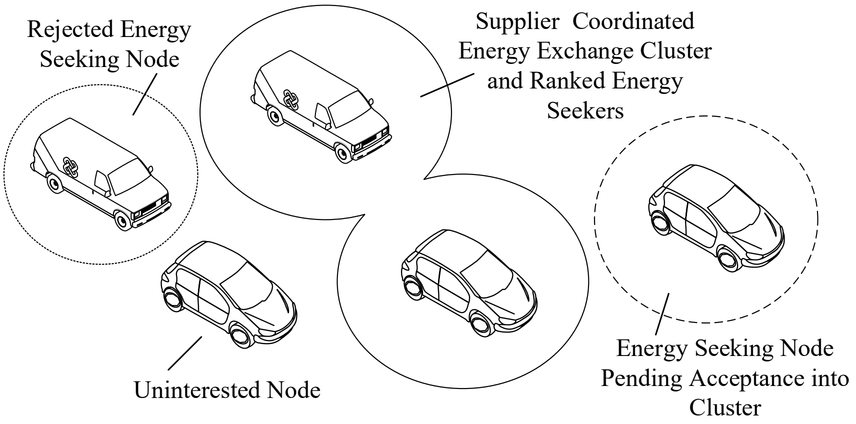
Rejected Energy Seeking Node
Supplier Coordinated Energy Exchange Cluster and Ranked Energy Seekers
Energy Seeking Node Pending Acceptance into Cluster
Uninterested Node

Figure 8 Supplier-Driven Energy Interest Clustering

Our proposed clustering protocol for CognitiveCharge nodes provide lightweight coordination mechanisms for contextual information dissemination and local energy exchange. These mechanisms are intended to operate atop, and be integrated with, any functionally compatible set of protocols. We therefore do not specify a precise format for messages, nor do we stipulate the underlying communication protocols. This section further highlights the scope of our cluster-based coordination mechanisms as well as presenting some key areas for consideration.

From the perspective of the cluster head, members of the energy interest cluster are considered part of a dynamic priority queue. Nodes in the queue are continually ranked by the supplier using our proposed CognitiveCharge Utility heuristic calculation in order to evaluate their suitability as candidates for energy exchange. Membership of an energy-interest cluster does not guarantee the opportunity for a seeker to acquire energy, nor does it guarantee an opportunity for a supplier to offload energy. Our supplier-driven approach dictates that the energy supplier decides which, if any, of the energy seeking nodes in the cluster it wishes to distribute amongst.

Cluster formation is initiated by an energy supplier via the broadcast of an ADVERT message to all nodes in direct communication range (effectively single-hop multicast to immediately connected nodes). Throughout the lifetime of the energy interest cluster, ADVERT messages are additionally sent from the cluster head to newly detected nodes which enter into single-hop communication range of the supplier. The advertisement messages signal to recipients that the source node has energy to offload, in addition to any details attached to the message pertaining to the supplier, such as the amount energy available for acquisition and the number of nodes actively queuing for energy.

Nodes with an interest in acquiring energy from the supplier indicate their desire to join the cluster by replying to an ADVERT message with a CONNECT message. To reduce protocol overheads and limit coordination complexity, in our approach we do not necessitate that the cluster head send explicit acknowledgement messages in response to requests to join the cluster. A join message sent to a supplier may contain additional data pertaining to the energy seeking node. As previously mentioned, joining an energy interest cluster does not guarantee that a node will receive energy, nor does it commit the energy seeker to acquire energy. The real-time exchange of energy between a supplier and a seeker is handled using the exchange protocol.

During the lifetime of the energy interest-based cluster, the supplier intermittently broadcasts UPDATE messages to cluster members in order to notify them of changes to the supplier's energy availability. These messages contain information regarding the ongoing availability of energy (e.g. the amount of energy left to supply). Due to the time criticality of energy exchange and the cross-layer geo-temporal dynamics of vehicular smart grids, it is important that the cluster head notify members promptly to changes so as to inform their own decision making (e.g. whether to leave the cluster). To allow for real-time responses to the evolving conditions, rather than sending UPDATE messages at regular time intervals they are sent reactively by the cluster head in response to change. The broadcast UPDATE messages may therefore additionally contain details pertaining to the priority supply queue, such as the length of the queue, duration of wait time, and position of nodes, etc.

Members can leave the cluster voluntarily (e.g. because they no longer need to acquire energy) by sending a DISCONNECT request. Similar to CONNECT messages, cluster members which choose to exit the cluster whilst it is active only need to send a DISCONNECT message if they remain in communication range of the cluster head. Moving out of range of energy exchange is interpreted by the cluster head as an implicit intention to leave the cluster. A cluster head can also evict a member from the cluster at any time using an EVICT message.

Peer-to-peer energy exchange is controlled via the supplying node, who, acting as cluster head, selects a cluster member – a self-reported energy seeking peer – for exchange using the CognitiveCharge decision making process. Exchange of energy between the cluster head supplier and a selected cluster member is initiated by the sending of SUPPLY message. Once energy exchange has begun, either the supplier or energy seeker may send a STOP message to indicate cessation of energy flow. As with data exchange and cluster membership, energy exchange may also be terminated via detection of a dropped connection.

An example of the complete process is illustrated in Figure 9, which shows a high-level overview of the energy exchange process when more than one seeker and supplier are in the same area and competing with one another for access to acquire (purchase) and offload (sell) energy. In Phase 1, Supplier_1 is connected to $\text{Seeker}_1$and broadcasts a single-hop advertisement message. This initialises an energy cluster comprising the two nodes. Needing energy, Seeker_1 requests to acquire energy from Supplier_1. This request is accepted, and Supplier_1 begins transferring energy to Seeker_1. In Phase 2, both Seeker_2 and Supplier_2 have come into the energy exchange range of Supplier_1 and one another. This extends the energy cluster to all four nodes. Both suppliers broadcast a single-hop advertisement message to the seekers in the cluster. Seeker_2 requests energy from Supplier_1, and this request is accepted. Supplier_1 ceases transferring energy to Seeker_1 and instead begins supplying energy to Seeker_2. In the final phase of the scenario, Seeker_2 ends the receipt of energy from Supplier_1. Note that this example omits the negotiation and transfer steps for simplicity.

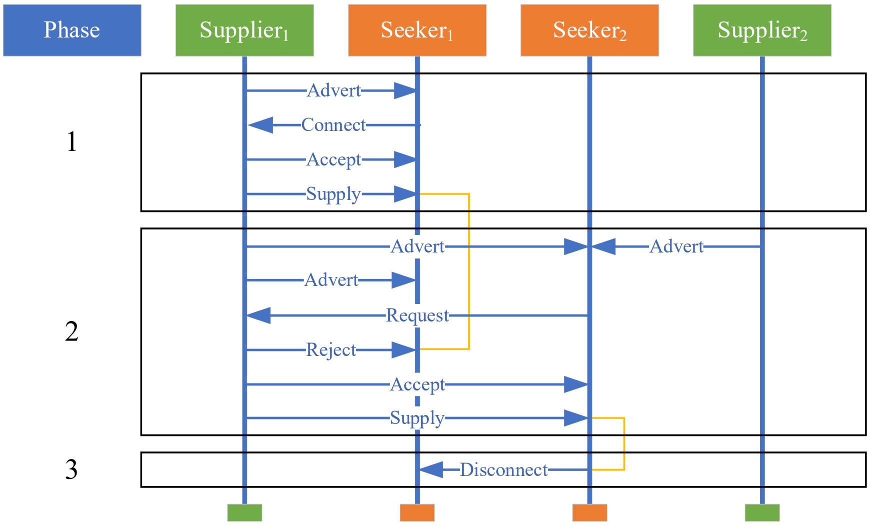

Figure 9 Energy Exchange Protocol Example

3.7 Analytics

3.7.1 Overview

Our predictive CognitiveCharge analytics are able to accurately anticipate future conditions by combining real-time contextual predictions with sensitive. Using these analytics, CognitiveCharge nodes are capable of decision making which avoids over – and under – reaction during periods of energy availability and threat environment instability. These analytics build on existing works for monitoring depletion, congestion, social dynamics, etc. in OppNets, in particular [38], [40], [42], [103]. A high-level overview of a single, generalised CognitiveCharge analytic is illustrated in Figure 10, which shows how each of our analytics combine multiple component metrics. Figure 10 further highlights how the metrics underpinning our CognitiveCharge analytics are captured asynchronously from multiple input signals. These inputs are observed from multiple external and internal dimensions as well as from second hand propagated information. Utilising established techniques for performant, real-time signal forecasting, raw data from first and second-hand observations are continually integrated into values which combine the anticipated value of a signal with an estimation of the signals level of volatility. This allows the raw estimated metric to be adjusted based upon the current degree of confidence in the prediction, given the recent context of the CognitiveCharge node. Each of the metrics (signals) underpinning our CognitiveCharge analytic can be considered as being updated at a regular interval.

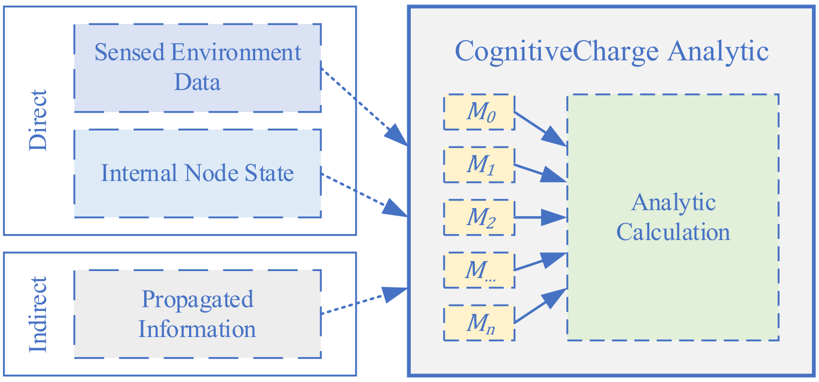

Figure 10 CognitiveCharge Analytic

Several utility methods are provided here which are used in calculating the individual analytics. For an input value x_t $(0 \leq x_t \leq 1)$ of a signal sequence X, the estimated next value x_{t+1} is calculated as s_t via the exponentially weighted moving average (EWMA) (Equation 7) using the smoothing coefficient α, per Equation 6. Whilst calculation of the EWMA provides an efficient method for forecasting, taken alone it does not capture the volatility of values in X. For instance, a stable, constant signal could be indiscernible from an unstable but periodic signal, reducing the reliability of forecast values. Therefore, to complement the value predicted via the EWMA, we also consider the degree of confidence in the estimation. The volatility of recently recorded values is measured as σ_t by the exponentially weighted moving standard deviation (EWMSTD) (Equation 8). Combining the EWMA with the EWMSTD for a given signal allows us to avoid overreacting to contextual volatility whilst increasing caution in decision making during periods of instability. A general formulation for the combined value estimation and volatility prediction for a metric is given as v_t in Equation 10 using the bounding function β in Equation 9. Finally, the method to rescale the bounded, volatility adjusted value to a new range whilst preserving the ratio is provided in Equation 11. This is used to ensure that each utility analytic is considered equally unless overridden by operator policy.

$$\alpha = \frac{2}{|X| + 1}$$

Equation 6 Smoothing Factor

$$s_t = \begin{cases} x_0, & t = 0 \\ \alpha x_t + (1 - \alpha) s_{t-1}, & t \geq 1 \end{cases}$$

Equation 7 Exponentially Weighted Moving Average

$$\sigma_t = \begin{cases} 0, & t = 0 \\ \sqrt{(1 - \alpha)(\sigma_{t-1} + \alpha(x_t - s_{t-1})^2)}, & t \geq 1 \end{cases}$$

Equation 8 Exponentially Weighted Moving Standard Deviation

$$\beta(x) = \begin{cases} 0, & x < 0 \\ 1, & x > 1 \\ x, & \text{otherwise} \end{cases}$$

Equation 9 Bounding Function

$$v_t = \beta(s_t \pm \sigma_t)$$

Equation 10 Volatility Estimate

$$\rho(v) = \frac{(v - r_{\min}) \cdot (r'_{\max} - r'_{\min})}{(r_{\max} - r_{\min}) + r'_{\min}}$$

Equation 11 Value Range Rescaling

The variable component metrics underpinning our CognitiveCharge analytics are continuously updated, triggered by events and interval expiration. This prevents nodes having to store large amounts of time series data and amortises the computational cost of calculating analytics over time, improving overall efficiency, and increasing the responsiveness of real-time decision making. The equations outlined are therefore incrementally updated and recalculated upon reading a new value x_t for a given input. As such, for implementation purposes, they are reformulated as iteratively updated functions.

The following subsections detail our proposed CognitiveCharge analytics. Each local analytic is explained in turn – first motivating the analytic, then describing the component metrics and providing formulations. After describing each analytic, we then detail our CognitiveCharge ego network analytics which extend the scope of the proposed local analytics to a node's socio-spatiotemporal community.

3.7.2 Depletion Rate

The heterogeneous nodes which form the VSG grid consume energy at uneven rates in order to provide services ranging from mobility (e.g. electric taxicab vehicles physically transporting people) to energy distribution (e.g. CSs selling energy on demand to neighbouring nodes). Energy consumption is nonuniform as expenditure of energy on services provided by nodes fluctuates due to the internal and external factors which influence energy availability and service demand. For example, electrical grid disruptions caused by fault or malicious behaviour can result in a significant reduction of local energy availability in conjunction with spikes in local service demand. Higher rates of energy expenditure from internal battery storage to provide services consequently increase dependency on local energy exchange opportunities in order to sustain service provision.

Our predictive Depletion Rate (DR) analytic accurately captures the complex dynamic relationship between multiple competing energy consuming services whilst additionally considering operator service provision policy (discussed later in this chapter). Rates of change in energy consumption per provided service are combined with predicted future capacity weighted by the dynamic service priority. This allows for nodes to predict the future energy availability requirements of themselves and others to a high degree of precision. By monitoring the outflow of energy per variably prioritised service, our Depletion Rate analytic allows CognitiveCharge nodes to balance adapting the dynamic priority of providing services with the necessity to acquire energy, given available opportunities for acquisition. Weighting of each energy providing service is defined by a set of externally defined functions assigned by the operator of the node.

CognitiveCharge nodes continually monitor the energy consumption levels of every energy-powered service they provide. For each provided service $s \in S$, the predicted level of energy consumption for each service can be represented as a sequence of values L_s. The DR for a single service is given in Equation 12, which shows how the EWMA and EMSTD use the first order difference in service level energy consumption. To combine the DR for all services, the DR is given as the weighted sum of the DR for each service, where each service weight is a value in $[0,1]$. A higher DR value indicates a worse rate of depletion (i.e. a greater degree of consumption and instability) and therefore a greater need to conserve or acquire energy.

$$
\begin{aligned}
L' &= \text{map}\left(l \mapsto \frac{l}{c}\right) L_s \\
L'' &= (\, l_i - l_{i+1} \mid l \in L' \,) \\
\text{DR}_s &= \nu(L'') \\
\text{DR} &= \frac{\sum_{s \in S} \text{DR}_s \cdot w_s}{\sum_{s \in S} w_s}
\end{aligned}
$$

Equation 12 Service Depletion Rate

3.7.3 Congestion Rate

The congestion rate (CR) of a CognitiveCharge node refers to the rate at which the queue for energy from that node is increasing. This builds on the CR as applied in networking [107]. As with DR, CR captures this information over time in order to anticipate and avoid cases of high congestion. This is advantageous not only from a time standpoint but from a service delivery perspective as well. Time spent unnecessarily waiting for an opportunity to acquire energy from a host is time that could be utilised to provide core services. For an EV, this would be mobility in order to conduct some activity. As an example, consider a taxicab awaiting the opportunity to charge is unable to transport customers between. The CR analytic allows for nodes to implicitly ascertain better opportunities for energy acquisition which reduce the impact of the charge event on service delivery.

A given node is limited in the number of peers it can simultaneously exchange energy with (e.g. a CS has only a limited number of outlets). However nodes can ‘queue’ for energy from both static and mobile nodes whilst others engage in transfer. By identifying nodes with high congestion rates, a node in need of energy can better identify underutilized nodes with surplus from which it can charge both from immediately and at future opportunities. As with DR, CR is a service specific measure as policies can determine that different nodes are treated differently. For instance, parallel priority queues depending on a node's affiliation. Under such cases, the CR for one service may be very high compared to another service for the same host node. The emphasis of CR is on capturing the immediate and anticipated length of the queue (Q) for energy. This is easily monitored as our CognitiveCharge clustering mechanism necessitates joining of implicit supplier-driven queues when nodes seek access to energy.

$$Q' = \mathrm{map}\left(q \mapsto \frac{q}{\max Q}\right) Q$$

$$Q'' = (\, q_i - q_{i+1} \mid q \in Q' \,)$$

$$\mathrm{CR}_s = v(Q'')$$

$$\mathrm{CR} = \frac{\sum_{s \in S} \mathrm{RET}_s \cdot w_s}{\sum_{s \in S} w_s}$$

Equation 13 Congestion Rate

3.7.4 Retentiveness

The retentiveness of a SG node refers to its ability to maintain charge beyond that which it requires for itself. This builds upon the data-context retentiveness analytic in [107]. Retentiveness (RET) is distinct from DR because RET focuses on the withholding of surplus energy whereas DR captures the actual expenditure of energy by services such as mobility. In other words, retentiveness measures how well a node can hold on to the extra energy it has. This concept is important because even if a node depletes its charge rapidly, it may still have a high level of retentiveness if it only needs to use a fraction of its charge for movement. Therefore, RET a crucial role in determining the efficiency and effectiveness of a VSG node's energy usage. As with CR and DR, retentiveness is dependent on the service context. Retentiveness is calculated based on the excess capacity at the time of each energy acquisition and the relative expenditure. In Equation 14, C is therefore the proportion of on-board energy available to expend by the node at the outset of each charge event.

$$\text{RET} = v(C\ )$$

Equation 14 Retentiveness

3.7.5 Receptiveness

Our proposed receptiveness (REC) analytic describes the predicted delay for a CognitiveCharge node until its next opportunity for energy acquisition from a suitable and willing supplier, e.g. charging from a dedicated supply station or available electric vehicle offering its surplus energy. The concept of receptiveness for CognitiveCharge nodes builds upon adaptive techniques which consider data storage space in forwarding decision making. Our predictive receptiveness analytic extends previous works to support capture the context and socio-spatio-temporal availability of energy exchange opportunities. This is accomplished by considering the delays between recent charge events (D) and using this to anticipate future charge opportunities (Equation 15), This complements our DR, CR and RET analytics.

$$\text{REC} = v(T)$$

Equation 15 Receptiveness

3.7.6 Incentives

Whilst not a fundamental analytic, incentive mechanisms such as dynamic pricing energy models can be integrated into CognitiveCharge node decision making. Incentives effectively enforce the CognitiveCharge utility-driven decision making. The price of a given neighbour’s energy is dynamic and related to implicit dynamically changing parameters which can be monitored and predicted in real time, as is the case for receptiveness. Dynamic pricing may be considered as a suitable real-world incentive for maintaining energy levels beyond need as well as permitting nodes to enforce a different preference regarding with whom, when, and how they choose to share energy. For example, a node may be well suited to providing energy but can negotiate through price to prefer providing energy to node it is friendly with, even if that node is in a less urgent energy state. Additional criteria for pricing based on cost of acquisition from diverse CSs can be considered. In [108] we considered a pricing formula which links directly to demand.

3.7.7 Trust and Peer Testing

Our lightweight CognitiveCharge proactive peer testing and trust mechanisms builds on existing works, adapted for the VSG context [70], [71]. They are designed to collaboratively evaluate, and disseminate information regarding, the energy behaviour of independent nodes in a VSG scenario. The testing is designed to ensure that nodes are behaving as expected, exchanging energy appropriately, and following established protocols for security and trust. By proactively testing the behaviour of nodes, CognitiveCharge can better identify misbehaviour. The risk associated with energy exchange in VSGs is context dependent and in some situations peer testing can be avoided; for example, in situations where there is a high degree of regional trust.The mechanism works by attaching test data to UPDATE messages to hide the intent and avoid test recognition. This technique seeks to prevent attackers from temporarily avoiding nodes that participate in collaborative peer testing and then resuming the attack once they have relocated. Figure 11 illustrates the peer testing mechanism. Node a first advertises itself as a supplier with small a surplus and receives a request for energy from node c. Node a sends an UPDATE message to node b with information about the request from node c. Node a supplies node c with the agreed upon amount and they subsequently disconnect. Node b then advertises itself as a supplier and receives a request from node c. Node b communicates via an UPDATE message to node a the details of the request. Where this is a discrepancy in the original and new amounts of energy asked for, nodes a and b are able to detect the misbehavciour of c and and mark it as distrusted.

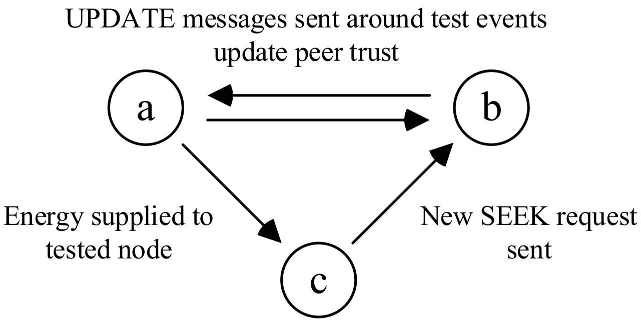
UPDATE messages sent around test events update peer trust
a
b
Energy supplied to tested node
New SEEK request sent
c

Figure 11 Peer Testing Mechanism

CognitiveCharge nodes prioritise neighbours with established feedback versus nodes with only a limited number of extremely high or low ratings. The value calculated for each node is weighted by the lower bound of the Wilson score confidence interval [109] [110], given as $\mathrm{L}_{95\%}(n)$. Weighting the overall utility of each node by $\mathrm{L}_{95\%}(n)$ prevents nodes with few ratings from skewing the selection of an available seeker or supplier against nodes with a slightly lower overall trust rating but a significantly higher number of reviews. For example, given 2 nodes, n_1 and n_2, with identical overall CognitiveCharge utility but the node n_1 having 1 positive review of 1 total and n_2 having 9 positive reviews of 10 total, node n_2 will be selected over n_1. This increases the resilience of the reputation system to the submission of false reputations by malicious nodes which measures such as the average or sum would be more vulnerable to. Trust and reputation values are considered reliable by default directly informed by binary ratings from peer testing.

3.8 Service Policies

In addition to the fully adaptive CognitiveCharge analytics defined in the previous section, it is also important to consider manually specified operator policies for energy service provision and security which can further alter behaviour. These were factored into our analytics as service weighting and influence adaptive decision making via weighting of our CognitiveCharge utility heuristic under certain contexts. It is important to note that these can be dynamic functions rather than explicit values. The following figure show an overview of the real-time CognitiveCharge decision making process, highlighting how externally defined operator policies are factored into the utility heuristic calculation. In this section we show how two generalised policy examples (for security and energy provision) are factored into the CognitiveCharge decision making process however any set of context dependent weighting functions can be specified by an operator.

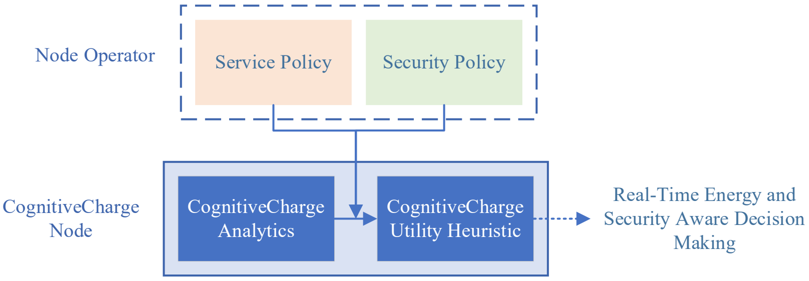
Node Operator
Service Policy
Security Policy
CognitiveCharge Node
CognitiveCharge Analytics
CognitiveCharge Utility Heuristic
Real-Time Energy and Security Aware Decision Making

Figure 12 Service Policy Integration into CognitiveCharge Decision Making

CognitiveCharge nodes can have different levels of service priority depending on their operators' preferences. A CS connected to grid infrastructure will likely seek to consistently maximise the amount of energy it offloads; conversely, a consumer electric vehicle will seek to balance energy offloading with consumption of energy to provide mobility. CognitiveCharge nodes are capable of supporting external operator preferences via weighting policies, which further improve prediction accuracy and inform real-time decision making. Dynamic service provision priority weighting policies specified by the operator of a node give precedence to behaviours under varying conditions. The simplest example for a set of services is to assign equal, uniform priority weighting. For example, the weighting function $\omega(s) = 1$ gives service s consistent, identical priority at any time. A more complex example for consumer electric vehicles would be to constantly prioritise mobility but take advantage of energy offloading opportunities to offset energy costs. This could be accomplished by assigning the energy offloading service a sigmoidal priority weighting function which reduces precedence of energy offloading behaviour when energy resources are actively, or predicted to become, constrained. E.g. $\omega(s) = \left(1 + e^{-10(\varepsilon - 0.5)}\right)^{-1}$. As with externally defined operator energy service weighting, CognitiveCharge also support operator policies for security thresholds. This allows for operators to specify adjustments to the level of trust required for autonomous participation in energy exchange. One such example is for preapproving certain peers.

3.9 Ego Network

Per [108], and building on previous works [40], CognitiveCharge additionally calculates and monitors ego network-based formulations of our proposed depletion rate (DR), congestion rate (CR), service demand (SER), receptiveness (REC), incentives (INC), and peer trust and reputation (TR) predictive analytics. These are prefixed with EN for clarity, i.e. EN-DR, EN-CR, EN-INC, EN-TR. Together these analytics describe the predictive multidimensional energy resource and security context heuristics of a node's ego network. Ego network analytics refer to the resource heuristics of the node's ego network. Ego network (EN) is defined here as a network consisting of a single node n together with the nodes that node n has encountered and gives each node their own perspective of the network. The exchange of ego-network analytics are done on aggregate via UPDATE messages which are reactively disseminated upon connection with peers.

CognitiveCharge allows nodes to aggregate resource observations disseminated by encountered nodes in order to form an ego network perspective of the network. Ego network information can be aggregated in many different ways and we have explored a number of models for weighting the contacts within a nodes ego-network in order to improve the accuracy of prediction of our EN analytics. Figure 13 shows how nodes exchange ego-network information via UPDATE messages and these are integrated into the separately monitored EN analytics. Through EN analytics, two nodes are able to share their perspective on the state of the network and further inform decision making.

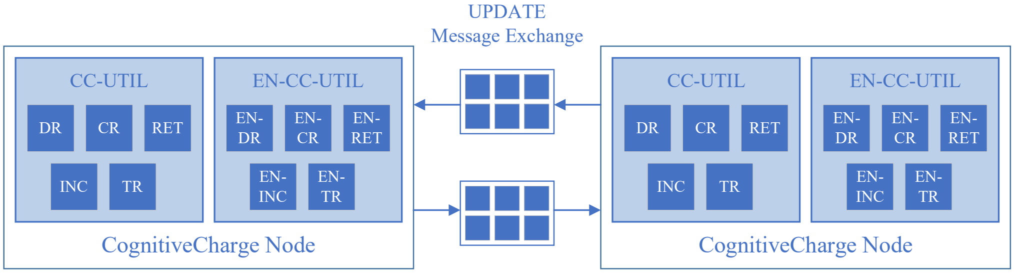

Figure 13 Ego Network Analytics Exchange

3.10 Decision Making

Our fully localised, real-time, heuristic-driven, cross-layer CognitiveCharge decision making process allows for nodes to adapt their service provision behaviour immediately in response to detected and predicted dynamic local and ego-network conditions. Our combined energy and security aware CognitiveCharge utility function extends previous works (e.g. [38], [42], [104], [107], [108]) with support for our refined and extended predictive analytics together, with additional support for operator service provision priority weighting policies. Conceptual state transitions, illustrated in Figure 14, are driven by our CognitiveCharge decision making process. The design of our analytics is such that there is no need for explicit thresholds and a value above 0.5 indicates a need to conserve energy and avoid exchange whereas a value below suggests that a node has energy available to exchange.

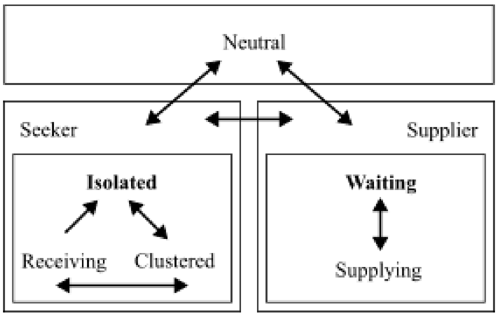

Figure 14 Decision Making State Machine

The decision as to whether or not to supply or charge from a given node at a certain point in time is based on the previously defined analytics. Where U is the set of utilities (calculated analytics) for each of the given criteria, we define the CognitiveCharge utility function for a given node n at time t as CC-UTIL (Equation 16). For a node actively seeking energy, over the set of potential connections which have a contact duration suitable for charging the highest (or lowest) will be selected for energy transfer. Higher and lower values capture the extreme of the need to acquire energy or the level of surplus.

$$\text{CC-UTIL}(n) = \sum_{u \in U} u(n)$$

$$\text{EN-CC-UTIL} = \sum_{u \in U\prime} u$$

$$\text{CC-UTIL}(n) = \text{CC-UTIL}(n) + \text{EN-CC-UTIL}$$

Equation 16 CognitiveCharge Utility Function

3.11 Summary

This chapter has described our CognitiveCharge proposal, which enables predictive energy availability discovery and threat context awareness to provide real-time identification and exploitation of suitably reliable opportunities for P2P energy exchanges. The architecture of CognitiveCharge combines data from multiple layers, including physical, data networking, and peer-to-peer energy exchange layers; social encounter dynamics layer; energy resource need and availability layer; trust and reputation-based security layer, and the supply and demand energy marketplace layer. The model of the VSG scenario is highly dynamic and multi-dimensional, and the ego network perspective considers the view of the network from the perspective of a single node. The chapter concluded by detailing CognitiveCharge analytics and decision-making.

Chapter 4

Methodology

4.1 Overview

In this work, we conduct simulation-based exploration of the performance of our proposed CognitiveCharge framework and protocol. Simulation-based experimentation is very widely used in computer network research [111]. Network simulators are frequently employed in the exploration and development of novel architectures and protocols as a time and cost effective alternative to prototypes and real-world deployments [112], [113]. Whilst technically practically possible in many cases, real-world deployment-based implementation and evaluation of computer networks requires extensive resources. As real-world networks continue to increase in scale in terms of the number of devices participating, so too must the scenarios we use to explore performance of new work. For this work, like many others, we find it infeasible to practically deploy a prototype with many real-world EVs spanning a suitably large geographic area and with the capability to exchange energy V2V. Test-beds such as MODiToNeS [114], whilst suitable for prototyping DTNs and OppNets, currently lack support for energy transfer and though reduced scale models could be developed, running many identical simulations with precise controls is impractical in the real-world. The VSG scenario central to this work is especially cost and time prohibitive to deploy to due to limitations in current, consumer hardware and firmware. Although a next step would be to be to look at small scale deployment of CognitiveCharge, in this work we focus on the initial feasibility and understanding of the performance characteristics of our proposal.

In designing and conducting the simulation scenarios we make use of a diverse array of heterogeneous VSG scenarios which have been selected to provide environments in which to explore in-depth the performance of our proposed system. Each of these scenarios is comprised of a range of consumer EVs in addition to static infrastructure CSs. The VSG scenarios fundamental to this research are complex, being both highly dynamic and multi-dimensional. The experimental scenarios used in our research seek to capture and represent these challenges appropriately. To conduct a comprehensive evaluation of the performance of our CognitiveCharge proposal, we must first implement support for the VSG, and subsequently a full implementation of CognitiveCharge, within a suitable simulator. The subsequent section of this chapter presents the mechanism for accomplishing this in greater detail.

For a thorough evaluation of cross-layer approaches to adaptive energy and security behaviours in VSG environments, it is necessary to suitably model each layer so that there is coherence within and across the simulated VSG layers. To augment the realism in our VSG scenarios and strike a balance across the diversity of our experiments, we make extensive use of complementary real-world data sets and pseudo-realistic algorithmic models to represent the fundamental facets of our experiments across data communications, energy exchange capabilities, mobility, energy resource-related activity, and malicious behaviour. For the performance evaluation of CognitiveCharge in this work, we utilise 3 distinct, multi-layer scenarios which capture the complexity of the highly dynamic and multi-faceted VSG. These scenarios were carefully selected to ensure that the evaluation covers a wide range of potential real-world urban and semi-urban use cases for VSG deployments in line with market penetration of EVs. By using these multi-layer scenarios, we can provide a comprehensive evaluation of CognitiveCharge and its ability to handle the demands of VSG in a real-world setting. The VSG scenarios used in our experiments are as follows: San Francisco, USA; North Somerset, UK; and a Manhattan Grid Model. A summary overview of the data sets and algorithms presented in this chapter and used for modelling aspects of our VSG scenarios is provided in the table below, which highlights where combinations of real-world data and complementary pseudo-realistic models are utilised for each layer.

Corresponding to the internal simulation scenario model, the remainder of this chapter builds up the simulation scenarios layer-by-layer. Firstly, the extended software simulation environment utilised for conducting experiments is detailed. Then, the specific physical characteristics of the VSG nodes that make up the experimental scenarios (the EVs and grid infrastructure CSs) are provided. The dynamic topologies of the scenarios, as driven by the mobility of EVs and which determine the network and energy connectivity, are subsequently detailed. Next is a comparative connectivity analysis of the VSG across a range of criteria. The mechanism for how fluctuating energy availability is modelled within each of our experiments and an analysis of the VSG scenarios is then presented. Formalising the behaviour outlined in Chapter 1 for the purposes of our VSG scenarios, the way malicious behaviour is incorporated into our experiments is then detailed. Finally, a summary of the VSG scenarios presented in this chapter is presented with concluding remarks.

4.2 Simulator

In alignment with the aims and objectives of this work, a network simulator is selected as the base simulator for this work due to the desire to apply opportunistic networking principles to energy exchange. In considering a network simulator suitable to model the VSG scenarios central to this work, we primarily considered established OppNet and DTN simulators in order to take advantage of existing software support within the simulator for data routing in these environments. Inline with our overarching aim, this would also facilitate integration and comparison with OppNet routing protocols as a comparison and benchmark. Whilst prior to this work no current simulator included support for VSGs, we consider a network simulator with support for DTNs and OppNets appropriate due to the additional needs of supporting network communications protocols and model wireless communications. Several network simulators are very established within the field, and so whilst these features could have been extended into, we found that it initially seemed appropriate to extend existing work and iteratively implement and integrate the necessary components to both model VSG scenarios and CognitiveCharge.

Many works have compared the performance and usability of different network simulators across different criteria (e.g., [111], [112], [113]). Whilst such works often highlight a wide disparity between the run-time computational cost and memory usage of network simulators [111], overwhelmingly, most are implemented as discrete event simulators [113]. A large number of network simulators, some previously very popular in research, are not actively maintained. For example, GloMoSim [115] (developed in C), ns-2 [116] (C++/Tcl), JiST / SWANS [87] (Java), and Adyton [117] (C++) are either inactive or have been superseded by more modern discrete event simulators such as ns-3 [118] and the ONE simulator [119]. Discrete event simulations are suitable generally favoured by the community as there is clear alignment between the event-driven nature of these systems and routing-related decision-making. We identified the ONE simulator as most suitable for this work due it's use as in OppNet research. The flexible and extensibile architecture of the Java source code lends itself well to the significant extensions required to fully implement and evaluate CognitiveCharge in multiple VSG scenarios. There is also a range of supporting literature and examples of significant extension by others (e.g., [120]). The architecture of the of the ONE simulator has also been used as a blueprint in the design of more recent simulation software packages [121].

Experiments in this thesis are conducted using the Opportunistic Networking Environment (ONE) simulator [119], with significant extensions developed through this work. The ONE simulator provides a comprehensive software suite for simulation-based evaluation of DTN and OppNet routing protocols and has been widely utilised in networking research. By default, the ONE simulator provides suitable tooling for in-depth exploration of routing protocols across various performance metrics in large-scale networking scenarios with variably configured devices. Using the ONE simulator it is possible to conduct large-scale, low-cost, reproducible experiments which would be infeasible using real-world deployments. However, complete support for the simulation of complex VSG environments is not currently provided by any single simulation environment. So as to adequately model the complexity of our VSG scenario, we significantly extend the core architecture of the ONE simulator with support for the additional layers necessary to realistically model heterogeneous EVs and CSs across the dimensions of the VSG.

Within the ONE simulator, nodes (effectively agents) are pre-configured with some initial parameters. For instance, a communication range and series of locations to move to within the simulation area. As nodes come into communication range, this presents the opportunity for nodes to exchange messages. An implemented decision-making process, typically a research routing protocol, then determines whether to create or send messages at any given time. Although events can be captured and played back to avoid simulating mobility, it is the core encounter events, usually driven at least initially by simulated connectivity and mobility (although there are real-world event traces in a format suitable for the ONE simulator), which are fundamental to the simulations and provide the opportunities for emergent simulated behaviour. In this work, complementary real-world data sets are combined with pseudo-realistic algorithms within the simulator in order to adequately capture the complex dynamics across and between layers of the VSG. In the ONE simulator, mobility is either given by events such as coordinates or driven algorithmically, such as a random waypoint mobility model. We implement support for overlaying additional, complementary data sets and behavioural algorithms so as to accurately represent the multiple attributes within and of each layer of the VSG scenarios. We do not abstract away components from our proposal in Chapter 3. Instead, we implement the full model, where support does not yet exist.

Figure 15 shows the original software architecture of the ONE simulator, highlighting how the movement models and event generators feed into the simulation engine. As data routing for DTNs and OppNets is the core focus of the ONE simulator, Figure 15 highlights how the configuration is used to derive connectivity during the running simulated scenarios. In a simulated scenario where mobility is defined (as opposed to the option to operate directly on connectivity events), nodes moving into communication range provide opportunities for communication which data routing protocols can exploit to transfer messages. We retain this core architecture in our extended simulation environment but inject support for energy connectivity and the associated internal decision making and data collection to facilitate this in implementing support for energy exchange decision making and, by extension, CognitiveCharge.

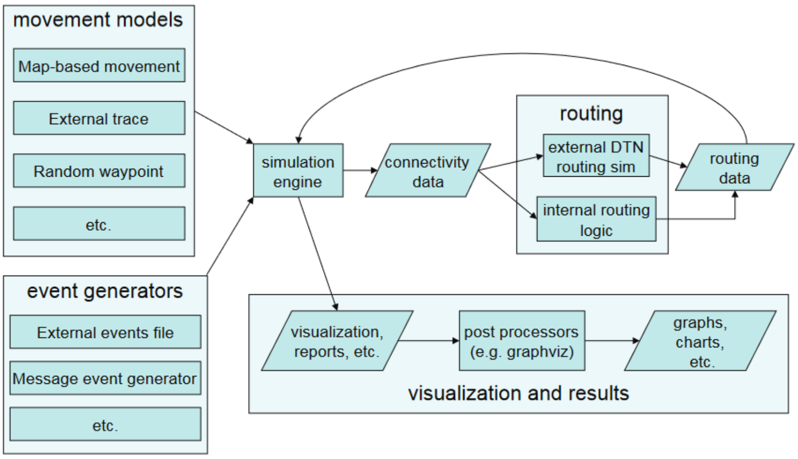

Figure 15 Architecture of the ONE Simulator [119]

The ONE simulator aims, where practical, to follow the typical networking stack. Applications operate strictly atop routing protocols. A provided example in the source code is the ping application, which is configurable to send ping messages to a range of nodes, which can be similarly configured to respond with pong replies. The routing of data messages through the network of nodes is handled by routing protocols, with messages passed over network interfaces which have physical properties such as communication range and link capacity. Figure 16 shows an example of a data message being sent from the application layer in Node A to Node B in the ONE simulator. The application is sending this message based on a simulation event. For instance, an automatic response to receipt of a message or a pre-scheduled time-bound message sending event. The message flow is highlighted alongside the object components that are required to be accessed in order to accomplish this. Conceptually, the message data (marked by solid arrows) flows mostly through the expected network layers (application, routing, interface, connection). However, components need supplemental information to be exchanged programmatically and access to shared objects to accomplish this (highlighted by the dashed lines).

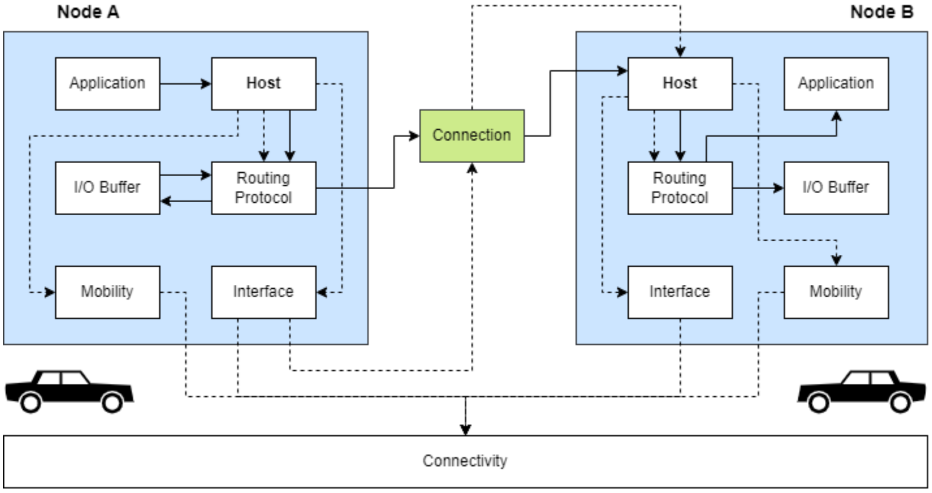

Figure 16 Message Sending in the ONE Simulator

A key challenge for extensibility in the ONE simulator lies in the way in which the simulation software maps partially, but not fully, to comparable real-world networking layers. This is notably distinct from simulators such as ns-3, where the simulation design far more closely follows the real-world implementation and system network stack. This blurring of network model design and simulation detail is part of what makes the ONE simulator a good platform for the rapid development and prototyping of new network protocols and suitably extensible for this thesis. Nevertheless, this raises conceptual challenges in handling states correctly, which can be complex and span components that might typically be considered out-of-scope. As an example, the underlying connectivity is optimised via a shared connectivity grid that is owned by an interface type. This allows for efficiently determining whether any pair of nodes are in communication range but blurs the line between simulation and model. Similarly, the 'host' or 'node' is considered to be the central hub to which all modules are attached. Some methods will directly call or programmatically address a host, whereas we would not have such reach in real models.

In order to add support for energy transfer in addition to data communications, we extend the ONE simulator with energy exchange functionality parallel to that of the networking components highlighted in Figure 16. A new CognitiveCharge decision-making component handles all coordination and decision-making between the two sides. Figure 17 shows a top-down view of the extended process of coordinating the exchange of energy from Node A to Node B, including message passing and energy flow between internal simulation software components (solid arrows) and object component access (dashed arrows). In Figure 17, the newly added components are shown in yellow, the modified components in orange, and the new CognitiveCharge decision-making engine is highlighted in green.

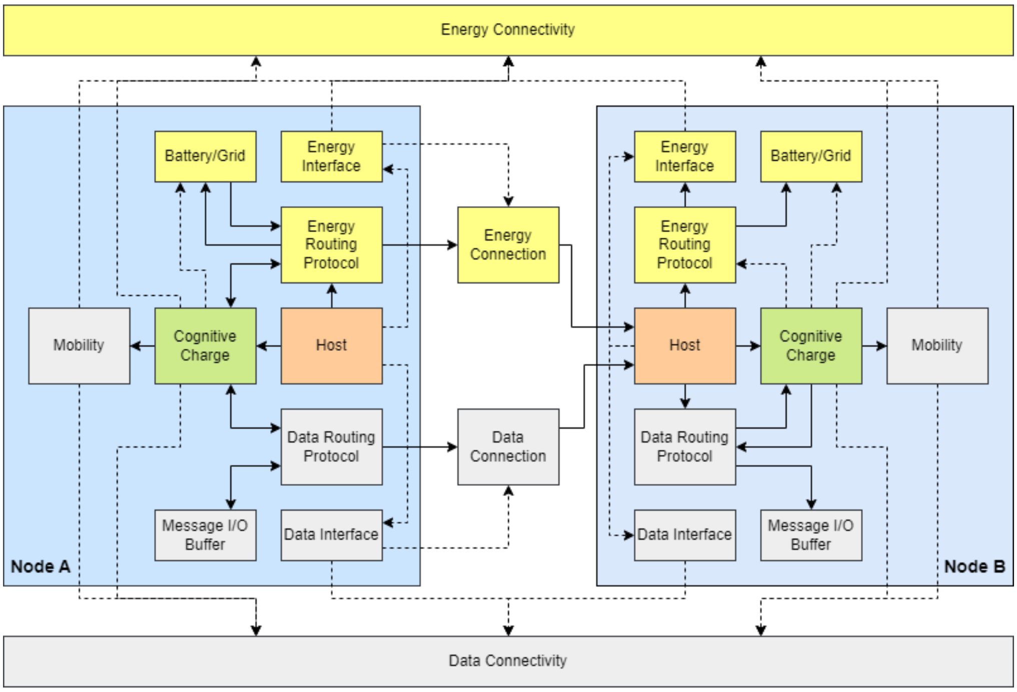

Figure 17 Message and Energy Transfer in the ONE Simulator

The new CognitiveCharge component is substantial and is fully responsible for host behaviour, together with the energy and data routing protocol components. This is a significant extension to the original mechanisms governing node behaviour in the default implementation of the ONE simulator. For instance, if a node representing an EV battery is depleted, the behaviour of the node should be that it cannot move. Similarly, if it is depleting and does not foresee a suitable opportunity for acquiring energy V2V, then it should relocate to acquire energy G2V from a CS. Such functionality was not considered in the core design of the ONE simulator, and so lots of objects and components needed to be linked via new behavioural controllers which were added to facilitate this. Metrics, such as battery level and contact durations, operate within their respective components, with analytics bridging between components to perform calculations atop these. The decision-making engine serves as the central hub and controller for all processes in the remodelled simulator and sits alongside the core host object, driving host behaviour from events within the simulation. In practical terms, the extensions made to the ONE simulator follow the existing design-patterns where possible, although Chapter 6 highlights limitations of this approach.

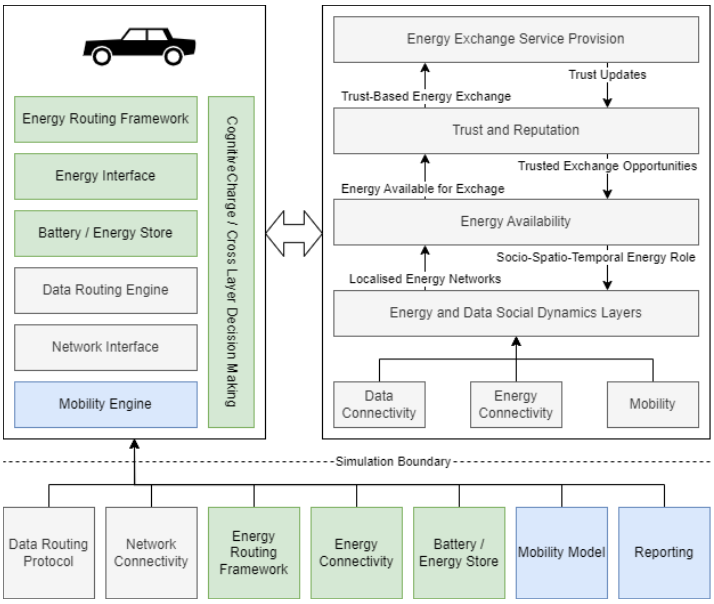

Figure 18 Physical Layers of the ONE Simulation Models

Figure 18 highlights how the physical simulation layers relating to the devices (EVs and CSs) are configured to drive the emergent simulated behaviour. Entirely new components are shown in green, modified components in blue, and original components in grey. By configuring the way each node should interact with each layer, the extensions to the ONE simulator, and the numerous additional programmatic alternations and new classes, allow us to model the full VSG as defined from the outset in Chapter 1.

The remainder of this chapter is arranged such that each layer of the internal simulation models of the VSG scenarios used in our experiments is detailed hierarchically in turn, from the physical layer at the lowest level, through to the energy exchange marketplace at the topmost level. Each of these layers represents a fundamental dimension of the VSG and is detailed in this chapter in the context of our simulation methodology.

4.3 Simulated VSG Devices

For real-time, fully localised data communications amongst future VSG devices, we consider all nodes as having mutually compatible networking technologies. Parameters selected for modelling data communications are chosen from the established literature [122], [123], [124] so as to realistically represent the 802.11p Wireless Access in Vehicular Environments (WAVE) standard whilst also suitably accounting for the constraints of non-line of sight communication between mobile nodes in built-up areas. As such, in our experimental scenarios, vehicular smart grid nodes communicate via Wi-Fi interfaces with a maximum range of 100 metres.

There are currently several competing standards for EV charging with varying degrees of interoperability amongst connectors. Likewise, there is a broad range of variably supported levels of charging speeds which differ across makes and models of EVs. Despite this, substantial efforts are being made worldwide by both industry and governments to establish common standards for EV charging. For example, the Combined Charging System (CCS) provides a set of open specifications which are seeing widespread adoption and have become a requirement of the European Union EV network [125], [126].

Of commercially available consumer EVs in the UK, over 93 per cent of models are already compatible (to some degree) with the CCS specification [127]. In line with the trend towards interoperable charging standards, in our experiments, we consider all participating EVs and infrastructure CSs as being universally compatible for exchanging energy. Furthermore, all EVs in our experiments seek to acquire and offload energy at their fastest supported charging speed.

As described in Chapter 3, we consider bidirectional V2V and V2G energy connectivity to be a subset of the network connectivity and require a direct P2P connection between nodes. Therefore, a given node cannot exchange energy with any other node if there is no direct, single-hop data connectivity between them to communicate information about the exchange. As we consider all EVs in our experiments to be communicating with one another via limited-range Wi-Fi interfaces, this means that EVs can only exchange energy with a subset of the nodes that they are immediately connected to and are therefore also geographically nearby.

EVs acquire electricity from the grid and store it in onboard batteries. This stored electrical energy is then consumed in order to provide services for end users. The primary service for which EVs expend electricity is physical mobility - the transportation of people and goods between geographic locations. For the VSG scenarios used in our experiments, the EVs can be categorised across several core characteristics; namely, the maximum available on-board energy storage, the rate of consumption in provision of services, and the rate at which energy can be acquired from peer suppliers in order to recharge batteries. In each of the VSG scenarios presented in this chapter, these fundamental characteristics of the EVs are modelled from real-world data. Figure 19 shows a scatter plot of the energy consumption, charge rate, and battery capacity of consumer EVs in the UK, categorised by the type of each vehicle [127]. Figure 20 highlights just the battery capacity of the EV categories.

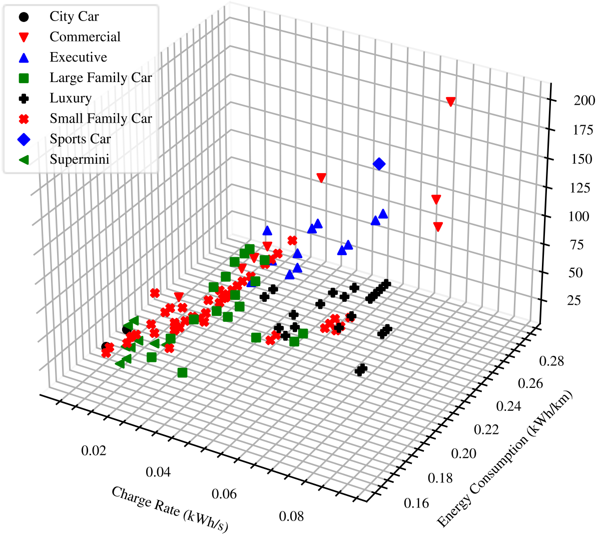

Figure 19 Categorised Electric Vehicle Data

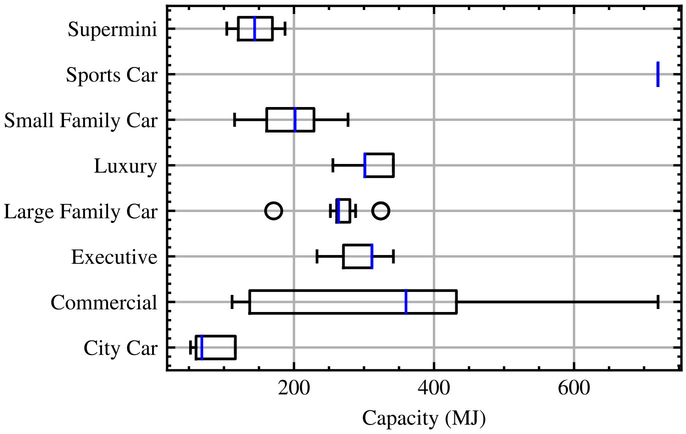

Figure 20 Electric Vehicle Battery Capacities

Values suitable for representing EVs in our experiments are derived from real-world data through taking the arithmetic mean of each of the relevant characteristics, further weighted by the number of vehicles for each category [127]. We exclude commercial vehicles and high-end/sports cars from the data for being unrepresentative of the majority consumer EVs. The resulting parameters generalise the EVs considered as VSG nodes in our experiments. An overview of the energy-specific configuration parameters for EVs used across our VSG scenarios is shown in the following table.

Table 2 Representative Electric Vehicle Data

Criteria	Value
Battery Capacity	200 MJ
Range	306 km
Consumption Rate	654 N
Charge Rate	50 kW

Conventionally, electric vehicles are marketed with energy capacity and consumption in terms of kilowatt-hours (typically denoted as kW·h or kWh) is a non-standard unit equivalent to 3.6 MJ. For consistency in this work, we use standardised metric SI (Système International) units. As an example, data from The Society of Motor Manufacturers and Traders (SMMT) shows that the top-selling EV in the UK in 2022 was the Tesla Model 3 [128]. The Tesla Model 3 is reported to have a 57.5 kWh hour battery, 250 miles of range, and a maximum energy exchange rate of 170kW [129]. Extrapolating and converting this data to SI units for our experiments gives a battery capacity of 207 MJ, a consumption rate of 514.5 N, and an energy exchange rate of 170 kW. Note that whilst this approach is suitable for the experiments in this work, there are important considerations of this model for future work. Further limitations of this approach are discussed in Chapter 6.

In early 2023, there were over 38,000 publicly available commercial CSs deployed across the UK, excluding private ones such those at homes and businesses [130]. A limited number of EV CS operators in the UK only support specific vehicle manufacturers through the use of proprietary connections and private owners networks. Overwhelmingly however, CSs are operating analogously to traditional petroleum refuelling stations and supporting all popular CS connectors and available to all customers via pre and post-paid options. In line with the current state of the charging infrastructure and with consideration of proposed and upcoming government regulations, in the VSG scenarios presented in this section all EV can make use of any CS which they encounter, providing that it is not already occupied by another user.

We assume that public CSs with two or more identical connectors at the same rate can charge the same number of cars concurrently at that transfer rate. Whilst many public CSs have multiple different connectors to support the various charging standards, the small subset of those which also support concurrent charging typically only do so at a reduced charging rate, with EVs sharing the available resources. CSs with multiple connectors on a single CS which all support the advertised charge rate, are sometimes referred to as truly simultaneous. EVs in our experiments seek to acquire energy from CSs at the fastest feasible charge rate. Accordingly, all infrastructure CSs dispense energy at the same rate. As stated in this chapter. The exact geographic deployment of static CSs in each VSG scenario various as a consequence of the way each is modelled and is explained in more detail in the following section.

4.4 Simulated Mobility

The core dataset for San Francisco contains mobility traces of 500 taxicabs in San Francisco, USA over a consecutive 30-day period. We consider the selected subset of 100 EVs in the real-world traces for San Francisco to be operating as Hackney carriages. In the UK, taxis are broadly licensed as either Hackney carriages or private hire vehicles. Hackney carriages operate independently and provide on-demand journeys to customers who can hail them curb side. This is exemplified by London's famous black cabs' which can pick up customers immediately when requested and take them to a requested destination. Conversely, private hire vehicles are required to be booked in advance of a journey and therefore make use of centralised scheduling systems to assign taxis to customers – for example the models used by Uber Technologies Incorporated and Lyft Incorporated which can provide booking via smartphone applications. Hackney carriages represent a suitable model for real-world consumer EV journeys as exact journey data is not established in advance. Unlike for centrally scheduled taxis, this prohibits precise, centrally coordinated planning of long-term recharging strategies and necessitates real-time adaptation to fluctuating energy availability and service demand.

The raw data for San Francisco contains a number of erroneous GPS coordinates. The presence of errors can be trivially confirmed as there are some coordinates recorded for taxicabs far into the Pacific Ocean, and some node speeds calculated from coordinate pairs exceed the feasible mobility of roadgoing vehicles. To keep a consistent strategy for erroneous coordinate removal, we use the speed and location of nodes for coordinate filtering. We set a threshold for legitimate values as those where the mobility of nodes does not exceed 77 mph (34.4 ms^{-1}), which is the maximum legal road speed in California plus a 10% margin of error to account for recording, map projection, and processing accuracy. Using this approach, we find a total error rate of 0.06%. These points are then omitted from the data by iterative removal with a total of 13 steps necessary San Francisco mobility data.

A choropleth map is shown in Figure 21, which highlights the percentage of logged coordinates in the dataset in San Francisco and each of the four neighbouring counties of Alameda, Contra Costa, Marin, San Mateo, and Santa Clara. The activity in San Francisco constitutes over 92% of the activity of taxicabs in the data that lies in the wider San Francisco region. Within the 117 districts of San Francisco itself, the Northeastern region sees the most taxi activity. The districts with activity levels above 5% are South of Market (11%), Potrero Hill (8%), Financial District (6%), and Downtown Union Square (5%). This can be seen in Figure 22 for the 121.5 km^2 are of area of San Francisco.

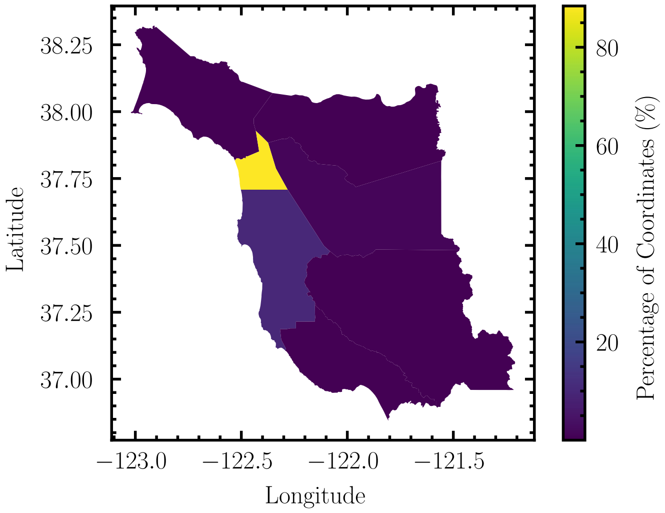

Figure 21 Percentage of GPS Points Recorded in San Francisco and Neighbouring Counties

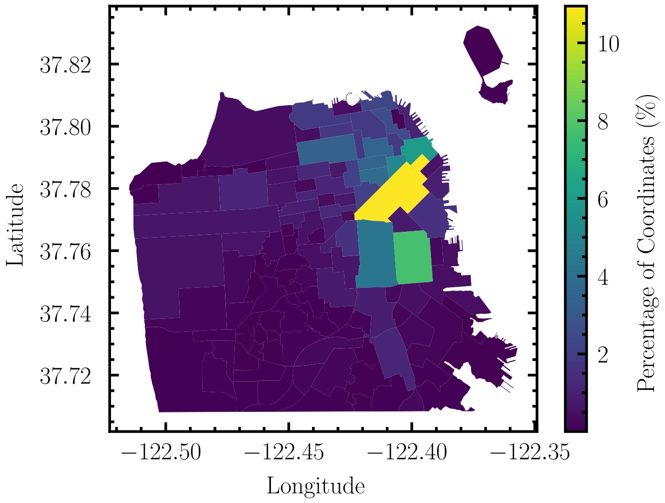

Figure 22 Percentage of GPS Points Recorded in San Francisco Districts

The ONE simulator performs poorly with very larges traces which require a high degree of accuracy and contain many nodes moving over extended periods of time. The combined data contains over 11,219,955 entries. There are 536 vehicles in the trace with a total trace duration of 84,330 seconds, resulting in a mobility trace containing 45,200,880 entries. Due to the extremely large simulation area, the memory usage of such simulations is very high because of the way the ONE simulator manages mobility and peer connection discovery. To overcome this, we extend the connectivity events supported by the simulator, encoding only meaningful trace data as events for the VSG scenario. As illustrated in Figure 23, each event comprises a tuple which contains additional spatiotemporal data. The data is processed once to obtain this data and then fed into the simulator. This is possible for real-world mobility traces as the movement of nodes does not change between simulation runs. A complete mobility trace is extracted from the processed data with points imputed at regular intervals. From this, an extended connectivity trace, as described, is calculated. For the limitations outlined, we take a 5 day subset of the data for 100 taxicabs where nodes are most active in the active region highlighted in Figure 22.

Because the mobility of nodes in the San Francisco scenario does not inherently support V2V or G2V energy exchange, we consider there to be an opportunity for exchange when vehicles are stationary for over 2 hours. This value was selected as it allows enough time for nodes to potentially charge from 0% to 100%. For energy exchange in the San Francisco VSG scenario, we are agnostic to the underlying technology but consider that nodes in 100 metre range can drive to one another in order to exchange energy. In simulations, we assume the centre of an energy exchange cluster in the San Francisco scenario to be the midpoint of all participating nodes. The amount of energy required to be consumed in order to travel to this point is then subtracted from nodes participating in exchange so as to account for the physical movement to reach energy exchange range. CSs are positioned where nodes are stationary for periods over 6 hours to account for end of shift EV charging.

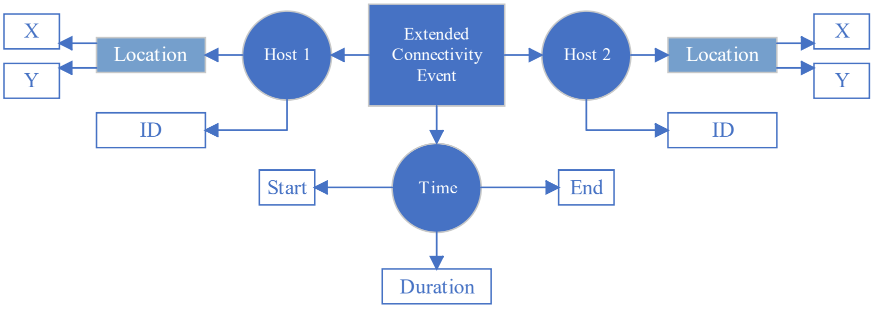

Figure 23 Extended Connection Event for Real-World Mobility Traces

North Somerset is a rural and largely agricultural district in the South West region of England with a population of approximately 216,000 and classified ‘Urban with Significant Rural’ by the Department for Environment

Food and Rural Affairs [131], [132]. Figure 24 shows the distribution of the population amongst settlements in the region, annotated with the number of settlements forming the population. Major towns are considered as having over 10,000 inhabitants, minor towns have between 1,000 and 10,000, and villages have under 1,000. The simulated road network for North Somerset has a total 1,329 km of drivable roads is extracted from publicly available data [133] and is shown in Figure 26.

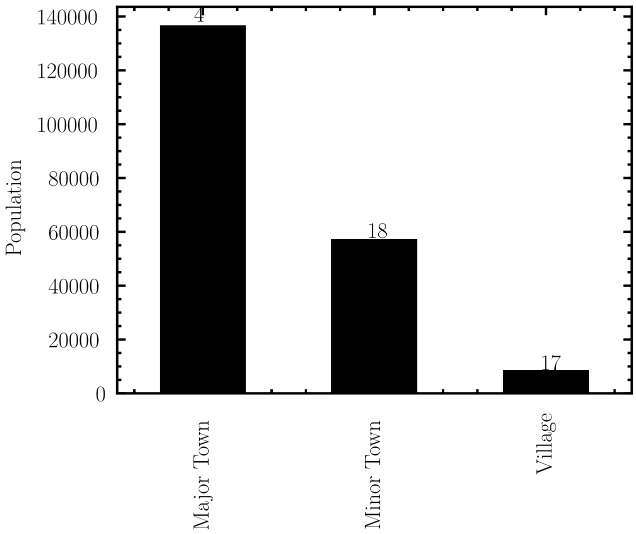

Figure 24 Settlements in North Somerset

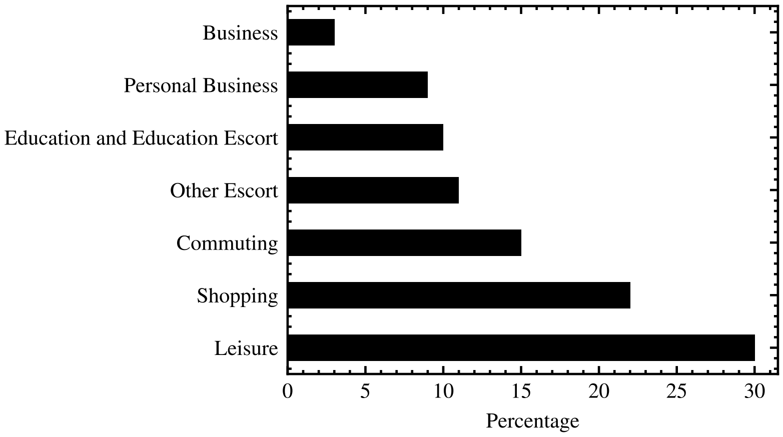

Figure 25 Categorised Vehicle Journeys

The dynamic topology of the VSG scenario for North Somerset uses a hybrid mobility model for VSG nodes, combining real-world data with pseudo-realistic algorithmic modelling. Movement of EVs within the North Somerset VSG scenario is determined according to a geo-social mobility model [134] which seeks to realistically capture the behaviour of drivers travelling with some regularity between commonly visited real-world destinations [133]. The geo-social mobility model sits in between purely synthetic mobility models, such as Brownian motion, and real-world mobility traces. This is accomplished by combining real-world locations of interest and social behaviour data together with a weighted, stochastic, destination selection and node activity algorithm.

CSs selected for inclusion in the scenario are distributed geographically according to the population of each settlement, where suitable CSs exist in the real-world. The ratio of EVs to dedicated CSs included in the North Somerset VSG scenario (1:200) is extrapolated from recent government data as the South West of England has 35 public CSs per 100,000 people and the national EV ownership rate is 7% [135], [136]. The North Somerset VSG scenario therefore comprises 200 EVs and 1 publicly accessible, dedicated CS. Additional CSs are available at the 50 included points of interest where these currently exist real-world. The major POIs are selected from real-world places, nominated geographically in proportion to the population density and categories shown in Figure 25 and Figure 24.

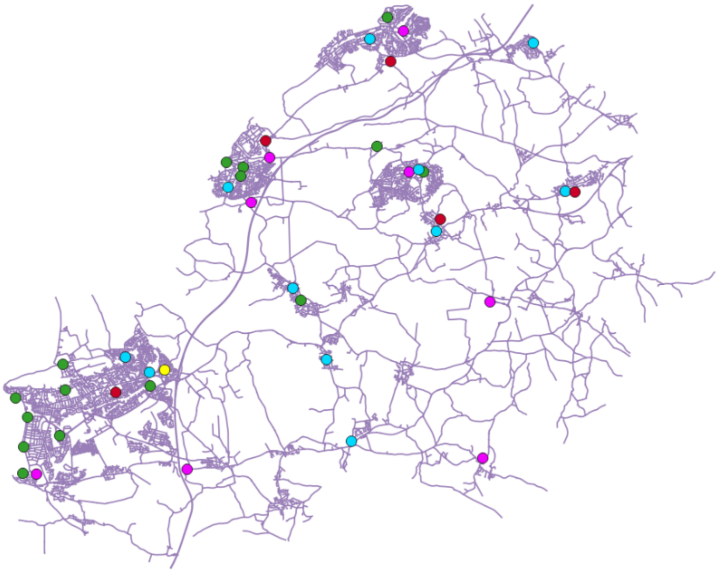

Figure 26 North Somerset VSG Map Highlighting POIs

Each EV in the VSG scenario has an individualised schedule of activities per day (e.g. a commute to a place of work), which will only be interrupted when a node has no alternative but to charge from a publicly accessible charging station. Figure 25 shows vehicle journeys made in the UK categorised by the journey type [137] which is used to provide journey information to the mobility model. The mechanism for determining when to disrupt a schedule is when the combined distance to the destination and from the destination to the charge point is infeasible, given the remaining battery capacity. The geo-social mobility model considers certain POIs compulsory and others more flexible. For instance, a commute to drop a child off at their school would be a persistent POI as the school would not change for the node. Conversely, a leisure trip would not be persistent as the same person may visit many different places for leisure (e.g. restaurants, gymnasiums, and cinemas).

The Manhattan grid model (MGM) is a pseudo-random VSG scenario. The purpose of including the Manhattan grid model in our experiments is to explore the effectiveness of CognitiveCharge in a scenario where EVs do not exhibit social behaviours. The Manhattan model VSG scenario in this work comprises 500 EVs and 50 publicly accessible CSs, based on the approximate ratio of EVs to CSs in Europe [138]. As shown in Figure 27, EVs move along roads in a square 25000 km^2 grid formed of square 500 m cells. The edges of each block represent roads along which EVs can travel. The combined length of roads in the scenario is consistent with the pseudorandom scenarios previously defined at 1,010 km. The velocity of EVs is between 10 and 15 ms^{-1} (approximately 20 to 30 mph) and the duration of the simulation is 5 days.

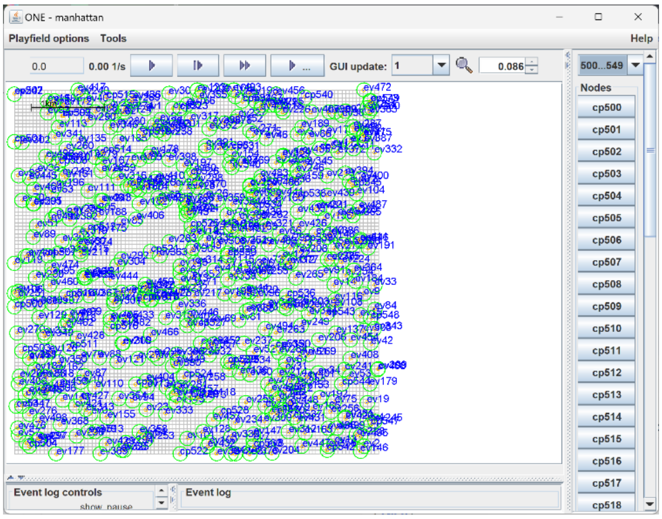

Figure 27 Manhattan Scenario Loaded in the Simulator

In our experiments, we employ the most widely used definition of the Manhattan grid mobility model (MGMM) wherein nodes can move along roads in a defined 2D grid area [139]. At each junction (road intersection), the probability of a given node making a left turn is 0.25, carrying on straight ahead is 0.50, and turning right is 0.25. These weights are scaled when reduced options are available, such as if a node can only go left or forward. U-turns are not considered possible in this model. In some versions of the MGMM, edge conditions are handled with nodes will wrapping around the simulation plane and re-entering the opposite side of the simulation area. To avoid this unrealistic behaviour, in this work, edges are considered as boundaries and nodes will continue moving with the reduced options. With probability 0.1, at each intersection, a node has the possibility of waiting for between 1,000 and 10,000 seconds before moving on to the next point. The wait time at all intersections is between 0 and 100 seconds.

We extend the MGMM for EVs such that nodes will visit a CS when the battery level reaches below 20% in order to avoid becoming stranded. As with the other VSG scenarios in this work, we presume that EVs have knowledge of the road map and locations of the static, infrastructure CSs. As such, nodes can determine with sufficient accuracy the energy cost of travelling to a CS versus the next junction. EVs in this scenario have a total battery capacity of 360 MJ. The CSs in our Manhattan grid VSG scenario are distributed randomly with the caveat that there is a minimum 2 km straight line distance between any two CSs.

This prevents unrealistic over clustering of CSs in any region of the grid. EVs in the scenario start in fully random positions on roads.

4.5 Simulated Device Behaviour

4.5.1 Energy Resource Behaviour

In all considered VSG scenarios, energy resource behaviour is governed by how a given node interprets and reacts its environment as well as by the primary behaviour of each node – mobility. For San Francisco VSG scenarios, movement of EVs is strictly governed by the real-world GPS coordinates and associated timestamps in the traces. For the North Somerset VSG scenarios EVs will move according to their core activity but seek to charge from an infrastructure CS when necessary.

In our experiments, we consider that the VSG infrastructure CSs behave consistently for all EVs in the VSG. Whilst it would be possible for CSs supplying energy to selectively accept or reject requests to charge from EVs based on arbitrary operator policies, e.g. vehicles of a certain manufacturer getting priority over others, in this work we consider their behaviour to be uniform. The CSs themselves will not display bias towards any EV. Likewise, the price of energy at each CS is uniform across the CS network and remains consistent throughout the duration of each of the scenarios. Our experiments therefore focus specifically on the behaviour of the decision-making of the roaming EVs. Variations on the approach concerning the behaviour of independent VSG nodes are discussed in Chapter 6.

4.5.2 Depletion Attack

In our VSG scenarios, we consider an active EDoS attack conducted by malicious nodes. Unlike with a typical energy theft attack, in an EDoS attack malicious nodes do not seek to only supply their own batteries. In order to maximise impact on the network, malicious nodes will continue to acquire energy beyond when their own batteries are full. Unlike a passive attack, in an active EDoS attack, malicious nodes will seek to promote themselves as being the best nodes to supply energy to. In these VSG scenarios, the attacking nodes use the false promise of payment as the means to promote favourably to suppliers. Therefore, whilst the EDoS attack is the primary focus of the malicious nodes, a side effect is the theft of energy as a result of the false promise to pay for energy acquired via V2V exchange. Compared to individual nodes engaging in energy theft, an EDoS attack is more devastating as, barring physical limitations, there is no upper limit on the amount of energy that can be lost.

In the experiments detailed in Chapter 5, malicious EDoS attackers were strategically placed at frequently visited POIs where infrastructure CSs were absent. For the North Somerset scenario these are the subset of the POIs defined. In the San Francisco we select POIs in the same way that locations of CSs were selected. Being geo-socially central within the network, these locations hold high significance. Attackers stationed here have a high degree of access to nodes in the network. As previously outlined, these POIs encompass a variety of sites, including shopping and recreational areas. From a networking perspective, these malicious nodes are effectively conducting an active black hole attack against the wider network except targeting the denial of access to energy instead of data.

4.6 Conclusion

This chapter detailed the experimental methodology we use in order to conduct extensive evaluation of our CognitiveCharge proposal. We utilise three diverse, multi-layer VSG scenarios which model the complex dynamics of the VSG: San Francisco, North Somerset, and a Manhattan grid model. Each scenario is multi-layered and combines pseudorealistic and pseudorandom data to suitably represent the complexity of VSGs. We extend significantly the ONE simulator, an agent-based discrete-event network simulator, for experiments in this work and implement in the extended simulator our CognitiveCharge framework and protocol. Initial configuration and scenario set-up is provided, with behaviour within the simulator derived from this. Further consideration of details of our experimental methodology are given in Chapter 6.

Chapter 5

Evaluation

5.1 Overview

This chapter presents evaluation of our CognitiveCharge proposal. Using the experimental methodology and VSG scenarios detailed in Chapter 4, we conduct a multi-criteria performance assessment across a range of measures to explore CognitiveCharge under a broad range of real-world, pseudo-realistic, and pseudo-random conditions. We first present the results considering only the exchange of energy in each VSG scenario. Secondly, we consider just the results pertaining to security. The final set of experiments bring together both the energy and security aspects.

Within the broader set of experiments outlined above, in this chapter, we group the results by the type of VSG scenario used to reflect the differences between them more clearly. The first group comprises the pseudo-realistic North Somerset VSG scenario wherein EVs will adapt their mobility according to the energy and threat contexts, as described in Chapter 4. The second group consists of the real-world San Francisco VSG scenario. This scenario is considered separately because EVs do not change their mobility based on the presence of any energy and security-aware decision-making processes; the mobility models in this VSG scenario are fixed. Finally, we consider the pseudo-random Manhattan grid VSG scenario as a separate group. Although EVs in the Manhattan grid will adapt their behaviour depending on the energy and threat contexts, this scenario comprises only public infrastructure CSs and EVs and vehicles move pseudo-randomly in an artificial environment.

5.2 Energy Exchange

In this section we compare CognitiveCharge against a range of decision-making approaches: grid-to-vehicle (G2V), rule-based decision making, and semi-centralised approaches. As a useful benchmark we first compare against a strict, unidirectional G2V energy exchange approach. This is most representative of EV charging in the real world. When using G2V decision making, CSs follow a strict first-in first-out charging policy. Under this policy, any EVs queuing for energy will be served according to the time they arrived. Breaks in connection with the CS, either via range-based disconnection or leaving the supply queue through the energy exchange protocol, will restart the connection when the connection resumes. Publicly accessible CSs order nodes by active connection time; this is analogous to the way most charge-points currently operate in the real-world today.

Lightweight, rule-based decision making involves application of rules to immediate circumstances without protracted data to provide context and inform real-time activity. We consider 3 rule-based approaches to V2V energy exchange as useful for comparison against our CognitiveCharge proposal. For threshold approaches to V2V energy exchange, we use thresholds set to 50%, 70%, and 90%. A threshold value of 100% is equivalent to a scenario where EVs will only acquire energy from CSs. We refer to these threshold approaches as T50, T70, and T90, respectively. For each of these methods, if the value remaining when the specified threshold is subtracted from the current energy level is positive, then an EV considers itself as having surplus and will advertise as being able to supply energy. As an example, an EV using the T50 decision making engine will seek to acquire energy if its own battery level is below 50%. When it has capacity above 50%, it will advertise itself as having a surplus of the current percentage subtracted by 50%. A high-level overview of the decision making for when a node is determining is own role based on threshold decision making is provided in Figure 28. We consider two peer selection mechanisms for threshold-based decision making. Namely, first-in first-out (FIFO) and last-in first-out (LIFO). These policies determine the order in which hosts will receive energy from the supplying node.

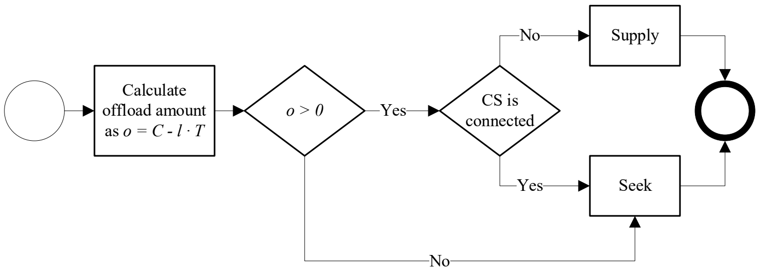

Figure 28 Threshold-Based Decision Making for V2V Energy Exchange

As explored in Chapter 2, the majority of approaches to P2P energy exchange in the literature are reliant on centralised or semi-centralised decision making. Therefore, we further CognitiveCharge compare against a cluster-based EV charging approach wherein a semi-centralised scheduler manages localised V2V energy exchange. Analogous to existing works, this approach uses a combination of pairwise ranking of EVs together with a dynamic threshold based on the energy levels of nodes in the cluster. EVs are paired such that nodes with above average energy in the cluster are paired for exchange with EVs that have below average energy. Pairs are arranged highest level to lowest level. Such an approach would not be possible under the stipulations we give for fully distributed and decentralised opportunistic VSGs (as detailed in Chapter 1), without lessening limitations on privacy. Under the supplier-driven and privacy-oriented behaviour of CognitiveCharge, information about a node's energy state is only selectively exchanged and nodes that wish to offload energy only need to indicate a surplus amount to potential peers. Nevertheless, the broader approach of centralised and semi-centralised decision provides a useful and informative comparison. We refer to this strategy as HL in the remainder of this Chapter.

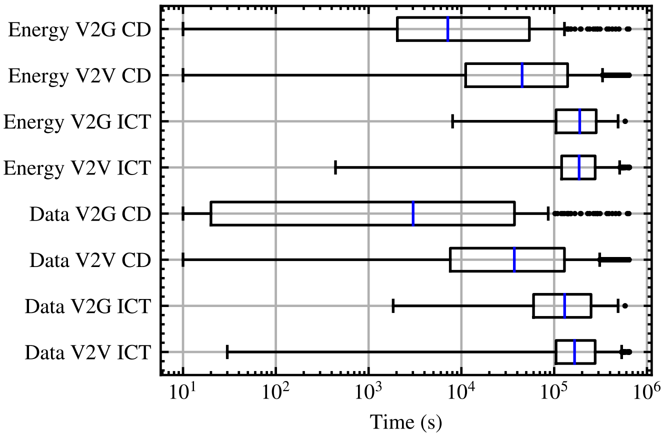

Figure 29 Data Connectivity Analytics for North Somerset

Figure 29 shows the contact duration (CD) and inter-contact time (ICT) connectivity analytics for the North Somerset VSG scenario under G2V charging. These are included pairwise for V2V connectivity and V2G connectivity for the EVs and CSs in the scenario. This is shown for both the data and energy connectivity layers. It is clear from these analytics that there is a significant amount time that nodes spend at, and in the vicinity of, CSs. As data communications have longer range than energy connectivity, we see the expected reduction in CD and ICT where nodes pass-by one another but do not spend significant time in the same vicinity. This is clarified in Figure 30 which shows how nodes will connect on the data and energy layer when in close range but only on the data layer when in longer communications range. These data layer connections are important for data dissemination; for CognitiveCharge this includes the exchange of UPDATE messages which contain ego network data and, when conducting peer testing, information about the recently tested exchange. The CD across the energy layer is typically longer because nodes are visiting specific POIs rather than just being in the same broad geographic region. Likewise, the ICT across the energy layer is longer as nodes are not frequently close enough for an energy layer connection. It is important to note that connectivity at the energy does not imply that two nodes are actively exchanging energy, only that it would be possible for them to do so should their respective decision-making processes lead to an energy exchange event. From the presented CD and ICT analytics, it is already possible to infer some behaviour. We know that EVs spend an average of 119 minutes at CSs but that the time to charge for each EV is approximately 67 minutes. This queueing highlights the necessity for our CognitiveCharge analytics (particularly here CR) and the need for improved access to CSs.

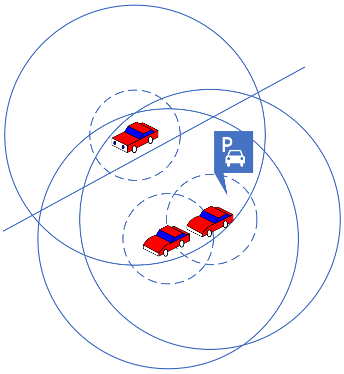

Figure 30 Range of Connectivity

An important consideration for the threshold approaches is that a practical implementation beyond the scope of these experiments is non-trivial. Such methods would be infeasible and would necessitate building upon a framework such as CognitiveCharge. In these experiments they operate atop our defined energy exchange possible. They are feasible in the simulated VSG scenarios because there remains a degree of homogeneity which does not exist in the real-world, outside of closely managed fleets. In reality, such approaches would necessitate significant additional information which would need to be supported by a framework such as CognitiveCharge. To be practically implemented, our CognitiveCharge Utility function (CC-UTIL) would need to be extended to support the threshold decision making mechanism.

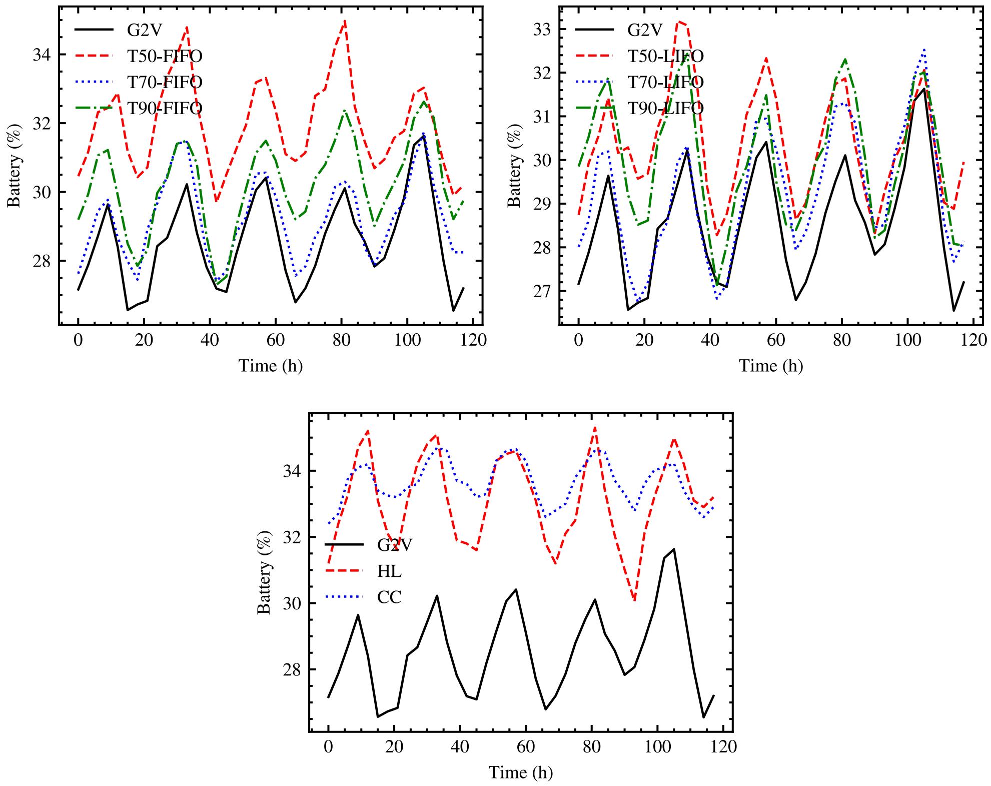

Figure 31 Mean EV Energy Levels in the North Somerset VSG Scenario

Figure 31 shows the mean energy level of EVs over time in the North Somerset VSG scenario for various peer-to-peer energy exchange strategies. From the average energy levels alone, the activity of EVs each day can be inferred. There is a clear cyclical pattern of energy usage each day with a general decrease during the day followed by an increase in the evening and overnight as vehicles are able to charge. The decrease in EV energy levels during the day corresponds directly with the consumption of energy through mobility which EVs expend as they carry out their behaviour. As EVs increasingly need to charge during the day, per the pseudo-realistic geo-social energy-aware mobility model utilised in the North Somerset VSG scenario, they consequently increasingly prioritise visiting a CS. We then observe the EVs acquiring energy from the CS before resuming the activity determined by the mobility model. Note that whilst the percentage differences in some cases seem numerically low, it is important to consider that just 1% of EV battery in these scenarios is capable of delivering approximately 2 km of range. Combined with the fact that there are large numbers of nodes in these scenarios, and it is clear how a 1% increase in mean energy levels could equate to an additional 400 km worth of additional mobility across the VSG.

The threshold strategies with LIFO queues perform similarly, with T50-LIFO, T70-LIFO, and T90-LIFO yielding an approximate increase of 1-2% across each day. The FIFO strategies show much greater variance compared to the corresponding LIFO approaches. T50-FIFO performs 4% better than the baseline G2V scenario, but T70-FIFO performs comparably to the scenario in which nodes may only acquire energy directly from the grid. Under both FIFO and LIFO threshold rule-based V2V energy exchange mechanisms, a threshold of 50 performs well as nodes continually seek to equalise the energy within their clusters. As can be expected, with the additional context available to the external decision-making engine, Figure 31 shows that even a simple semi-centralised approach performs well and shows a consistent increase in the average amount of energy EVs have available. On several occasions, HL increases the average energy slightly beyond CognitiveCharge, despite the lack of transferred knowledge. We can expect this to be the result of more rapid movement of energy amongst EVs. The semi-centralised scheduler is able to assign and adjust pairs of nodes so that there is a continual evening of energy within each cluster. However, whilst CognitiveCharge is unable to increase the average battery level beyond the HL strategy, it is clear that the energy levels of EVs remain more stable over the course of the simulations. As a result of our CognitiveCharge socio-spatio-temporal analytics and real-time utility-drive decision making, EVs have a much greater awareness of their own energy needs as well as the energy state of their ego networks. This allows for nodes to avoid blocking access to temporary EV suppliers and CSs to the nodes with the most critical need for energy, as better alternative opportunities for future access to energy can be identified. Although not possible in this work given our stringent definition of the VSG as necessitating localised, distributed, and decentralised decision making, we can see from the results how a fully centralised strategy which combines the protracted knowledge and network awareness of mechanisms such as CognitiveCharge with global scheduling could readily outperform all of the approaches deemed viable in our defined VSG.

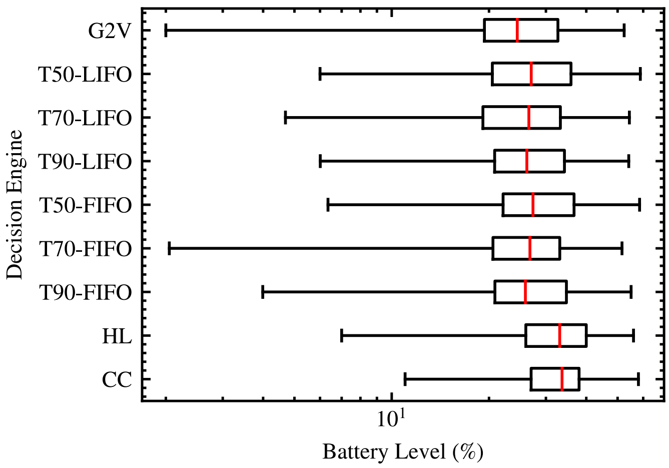

Figure 32 Distribution of EV Energy Levels in the North Somerset VSG

Whilst changes to the global average are helpful indicators to the general state of the EVs, they alone do not necessarily reflect an improvement for all nodes and at face value can be misleading. It is possible that whilst the average has increased, some nodes are significantly better or worse off at the expense of others. For additional insights into the effects of each decision-making engine on the overall energy levels of EVs in the North Somerset scenario, it is necessary to consider the EVs at either end of the energy spectrum. Figure 32 presents boxplots of the distribution of EV energy levels across the duration of the simulated VSG scenario. The amount of energy available to EVs in the scenario is limited by the amount of energy in the network at any particular time. The rate at which energy can enter the network is limited by the EVs acquiring energy infrastructure CSs. For instance, in a small network with 2 EVs and 1 CS, the maximum rate at which new energy can enter the network is limited by the time that either of the EVs spend acquiring energy from the CS. In each of the VSG scenarios, the CSs are in near constant use so the upper limit of energy is similar for each. For this reasoning, it is only possible for V2V energy exchange strategies nodes to have limited impact on the overall or upper end of EV energy levels. Where these strategies have greater effect is on the distribution of energy within the VSG scenario.

We can see from Figure 32 that the similar behaviour for FIFO and LIFO threshold strategies are echoed generally for the distribution. LIFO strategies perform similarly and clearly increase energy levels for the nodes at the lower end of the scale when compared to the baseline G2V energy exchange approach.

The FIFO approaches are more unpredictable, particularly for the nodes with the lowest energy levels. Altogether, this suggests that as simple strategies for V2V energy exchange can form the basis for limited knowledge approaches, with the caveats previously outlined. FIFO based approaches are more dependent on selection of a suitable threshold value, but accurate threshold selection can yield better results than LIFO approaches. We can observe that rule-based approaches generally yield similar, lower average energy levels across the scenario than the baseline G2V case. For the more successful methods, this is the result of the energy being distributed appropriately to lift the energy levels of the most vulnerable nodes. For HL, we can see that the increase in average energy levels compared to G2V is to the benefit of all nodes in the scenario. This is evidenced by the increase in the upper and lower quartiles as well as the mean. It is important to note that under the G2V, HL, and T50 energy exchange approaches, no nodes end up fully depleted of energy. The results shown for the North Somerset VSG scenario across the threshold decision making, semi-centralised approach, and CognitiveCharge, are representative of our findings in all three VSG scenarios.

As an indicator of the success to which our proposed CognitiveCharge analytics are capturing the intended VSG context and features of the VSG, we can explore correlations between relevant in-network features and the values of the analytics. Here these have been monitored independently of our CognitiveCharge decision-making mechanism as this would unintentionally influence them. For DR, we observe an extremely high correlation of 0.98 between the DR and a node's self-identified need to acquire energy under the G2V scenario. For threshold-based approaches, this drops to 0.52 as the urgency of the need for energy is mitigated somewhat by the decision-making engine. The drop in the correlations between analytics and scenario features across the baseline and threshold approaches is expected behaviour. This is because the unintelligent opportunistic nature of the rules governing energy exchange do not account for the nodes actual need and can affect nodes counter to their need, e.g. a node with surplus receiving even more. In the case of DR for example, this means that the accuracy of such metrics is less effective at capturing the need of nodes at middling energy levels. For the baseline scenario where EVs exchange with CS only, we see similarly high correlations for congestion rate

(CR) versus immediate queue length (0.86), retentiveness (RET) versus surplus (0.73), and REC versus charge interval delay (0.81). The usefulness of these correlations in isolation is limited, however they serve to indicate that our proposed analytics are capturing the intended information and informing decision making as designed. The collective effect of our CognitiveCharge decision making is demonstrated through the overall energy level increase.

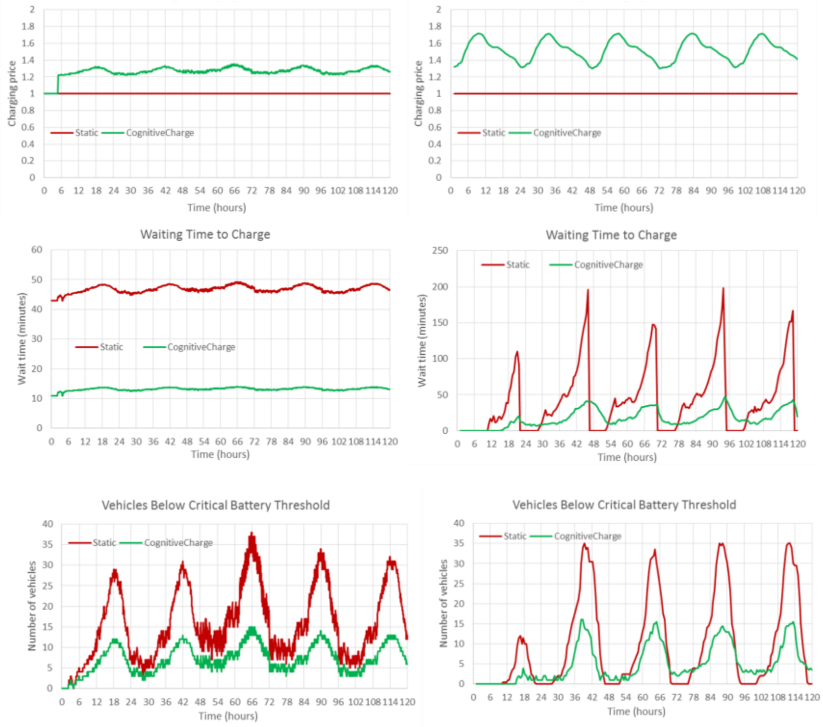

Figure 33 Impact of Incentives on Wait Time and Criticality Levels in San Francisco (Left) and Nottingham (Right)

An optional, although practically necessary, analytic integral to CognitiveCharge, motivates behaviour through incentives (INC). A relatively simple dynamic energy pricing incentive scheme can reduce both the wait time at CSs and reduce the number of nodes in need of energy. The experiments which derived the results in Figure 33 were conducted using traces of San Francisco (USA), and Nottingham (UK) [108]. The Nottingham Scenario was modelled very similarly to the North Somerset scenario. These figures have been reproduced here to illustrate how incentive schemes can play an important part in CognitiveCharge decision making. Nevertheless, as noted in Chapter 3, the economics of energy pricing and incentive schemes are beyond the scope of this work. The experiments presented in this Chapter consider CognitiveCharge nodes to be non-greedy and cooperative and so our CognitiveCharge utility function (CC-UTIL) alone determines behaviour. This means that all EVs in the VSG scenarios will all adhere to the internal decision making as to when to acquire or offload energy. Nevertheless, the pricing incentive mechanism for CognitiveCharge is necessary to enforce and regulate this behaviour in real-world deployments where independent nodes might act greedily or selfishly. A pricing strategy in such circumstances helps balances need and motivates the selling of energy when there is regional deficit but local surplus.

5.3 Depletion Attack

This section explores the performance of the CognitiveCharge peer-testing and threat detection mechanism for threat detection. In particular, the accuracy of attacker detection in each of the scenarios presented (Table 3) in the presence of increasing numbers of malicious nodes. We measure the detection accuracy of CognitiveCharge as the correct identification of a malicious EDoS attacker node, once tested, by peers who either directly tested the peer or have received an ego network exchange that includes the trust rating of the peer.

In the North Somerset VSG scenario, under active EDoS attack we find that CognitiveCharge has a consistent 100% accuracy in detecting malicious nodes. This is not a surprising result, as the confidence in the outcome of each test can be very high due to both the attacking and non-malicious nodes being stationary in the same vicinity for extended periods of time. There is no delay between tests and there are multiple peers available for testing due to the nature of the deployment of POIs. Whilst the variety of the nodes visiting each POI is low due to the geo-social mobility model, this has an advantage in quicker consensus reaching regarding identification of malicious nodes. Nodes with lower ties to these nodes are therefore more likely to receive a stronger consensus about the attacker. However, consensus is only useful for the nodes that will actually encounter the malicious attacker and a downside of this is that broader dissemination is limited. The results are identical for the Manhattan grid scenario as despite the reduced inherent social behaviour, the high node density yields the same effect.

Table 3 CognitiveCharge Attacker Detection Accuracy

	5%	10%	15%	20%
North Somerset	100%	100%	100%	100%
San Fracisco	68%	65%	57%	43%
Manhattan	100%	100%	100%	100%

In contrast to the North Somerset and Manhattan VSG scenarios, the detection accuracy for nodes in the San Francisco is far lower at 43-68%. This is largely due to the lack of peers available for immediate testing and the lower tie-strength. The local exchange clusters are far more sparsely connected and consequently there are also fewer events in the San Francisco scenario to conduct peer testing. Whilst being ostensibly an urban scenario, the energy connectivity is lower than the North Somerset scenario, which limits the broader applicability of this result. Nevertheless, we would intuitively expect results to lie somewhere closer to the North Somerset scenario in urban environments. This is particularly true given the more realistic energy seeking behaviour around CSs and POIs. It is clear, however, that the peer testing strategy loses accuracy when there is increased delays between data exchanges and an area for future work will be in further exploring and augmenting this approach in more sparse networks. It is worth noting that the limitations of the San Franciso scenario employed extend here as we observe that in the time period, where not immediately detected, the attackers are not later detected due to a lack of connectivity. There is insufficient ego-network data exchange in the limited time frame to disseminate information about attackers to future visitors. This is a well-known challenge of sparse networks and has unfortunate effects. Mainly, the lack of information dissemination beyond each nodes' ego network reduces the utility of the proposed strategy on more transient nodes with weaker socio-spatial ties. An broader concern is that attackers could more easily change their behaviour to opportunistically target specific nodes. Multi-hop reputation dissemination strategies have been widely explored in the literature and are an important consideration for future work in larger scale VSG scenarios where there would be many more overlapping geo-social networks. Whilst beyond the scope of this work, such an approach could complement CognitiveCharge and reduce attack impact whilst simultaneously increasing detection accuracy through greater information exchange. Nevertheless, such approaches come with their own challenges and are not a certain panacea.

5.4 Combined

The previous sections of this chapter have looked at the performance of a range of energy exchange protocols and the threat detection and mitigation component of CognitiveCharge separately. This final section builds upon these and looks at the complete picture of CognitiveCharge in VSG scenarios under active attack by malicious nodes.

As already noted, CognitiveCharge deployed in the North Somerset and Manhattan Scenarios has a 100% detection rate and rapid detection time due to the denser scenarios. In these VSG scenarios, the impact from attack is less than 1% in all cases and can attributed to the initial peer tests conducted in order to detect the attackers. The energy lost through these mechanisms is recovered quickly and there is no discernible impact beyond this. Conversely, the sparser San Francsico VSG scenario is more interesting and instructive to explore here. Figure 34 shows the energy lost to malicious attackers in the San Francisco VSG scenario when nodes are using our CognitiveCharge decision making framework and, for comparison, when nodes are using the semi-centralised HL strategy.

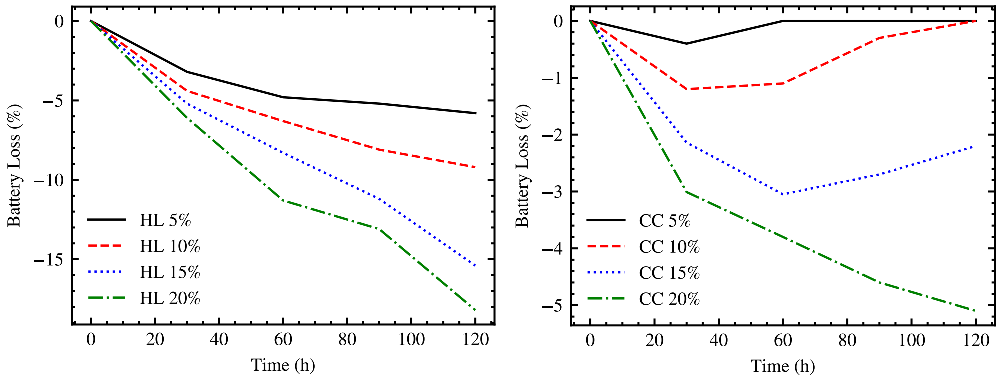

Figure 34 Energy Lost by EVs to Attack in the San Francisco VSG Scenario

From the outset it is clear that even an uncoordinated EDoS attack conducted by malicious nodes is a viable mechanism for disrupting energy flow in VSGs. As can be seen in Figure 34, the impact of increasing numbers of attacking nodes leads to significant loss of energy and the knock-on effects over time are exacerbated by continual energy loss. Whilst comparison against HL is a helpful indicator of the impact of EDoS attack in untrusted V2V energy exchange scenarios, it is nevertheless a somewhat contrived example. We include it here as worst-case scenario where an undetected attacker is involved in a significant number of exchanges due to the falsified messages positing it as being most in need. Nevertheless, the advantage of any degree of centralisation is the ability to coordinate and exchange information more rapidly so in reality only an unrealistically naïve version of HL would be impacted to this extent.

Figure 34 shows that CognitiveCharge nodes are able to detect malicious attack and that once detected, the impacts of the attack are reduced. Whilst Table 3 showed that detection rates were lower in the San Fracisco scenario, the impact of the attack is still reduced significantly for nodes using CognitiveCharge. It is clear from Figure 34 that even where malicious nodes are undetected, CognitiveCharge nodes can better manage the available energy resources and avoid offloading energy. As the impact of the attack and loss of in-network energy reduces local battery levels, CognitiveCharge nodes that previously had surplus shift their behaviour to conserve energy. At lower levels of malicious nodes (5% and 10%), this inherent behaviour adaptation fundamental to CognitiveCharge facilitates recovery of energy and restoration to previous energy levels in under 120 hours. The findings presented here and discussed in this Chapter demonstrate both the benefits of our collaborative peer testing mechanism but also the innate resilience of CognitiveCharge in untrusted VSG scenarios and in the presence of malicious nodes conducted EDoS attack. Not only can CognitiveCharge effectively detect threats, but it also mitigates the impacts of attacks through adaptive energy-resource management, even with large numbers of malicious nodes and when attackers remain undetected.

Chapter 6

Discussion

6.1 Overview

This chapter considers the wider context of CognitiveCharge. In this work, we focused on VSGs comprising EVs and infrastructure CSs. However, the concept of P2P energy exchange is more broadly applicable beyond this scenario and has been explored to varying degrees in the literature. This chapter looks at alternative application areas, such as unmanned aerial vehicles (UAVs) and unmanned underwater vehicles (UUVs). There are also a number of considerations to be made relating to the experimental methodology, it's limitations, the wider applicability of results, and the broader context of this work. This Chapter also considers these in more detail; providing discussion of limitations of current simulation software and further motivating the need for accurate and performant next generation simulation network environment tooling.

6.2 Assumptions

As detailed in previous chapters, this work presumes a core VSG scenario such that our proposed CognitiveCharge framework has the broadest applicability. By focusing on a stringently defined VSG use case, CognitiveCharge is highly extensible and adaptable to other areas where there is either overlap or there are less restrictive VSG requirements. Figure 35 shows the range of effectiveness of CognitiveCharge. This diagram seeks to highlight where CognitiveCharge can have benefit to similar SG environments given varying security demands and energy policies. For instance, in a scenario with maximal security requirements, untrusted nodes would simply be avoided altogether. Likewise, in a scenario wherein nodes will only conserve energy and cannot be inventivised to participate in P2P energy exchange, CognitiveCharge will have no benefit. Nevertheless, these situations could still benefit from CognitiveCharge where operator policy provides favoured, trusted peers.

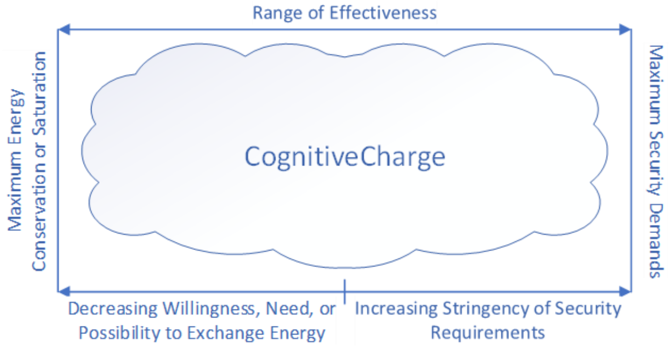
Range of Effectiveness
Maximum Energy Conservation or Saturation
CognitiveCharge
Maximum Security Demands
Decreasing Willingness, Need, or Possibility to Exchange Energy
Increasing Stringency of Security Requirements

Figure 35 Range of Effectiveness of CognitiveCharge

In this research, we operated under the presumption that Electric Vehicles (EVs) are uniquely associated with specific individuals, meaning that each EV can be linked directly to a particular owner or user, and their usage patterns, preferences, and behaviours could be traced back to that individual. This foundational assumption influenced the parameters and outcomes of our experiments. Vehicles remain at a particular POI whilst the owner or user is at the location. Whilst this is a reasonable assumption for EVs, and closely aligned to the real-world, increasing numbers of autonomous passenger vehicles may change this in the future. It is safe to assume that as the taxicabs that made up the original nodes were petroleum powered, they will not fully realistically represent EVs' energy-seeking behaviour. For instance, petroleum-powered vehicles are able to refuel at dedicated stations within just a few minutes which is not possible for EVs today.

The VSG scenario outlined in Chapter 1 assumes a maximally privacy-oriented scenario with minimal information visibility. Increased information dissemination, such as the details of energy exchanges, is to the benefit of more informed decision-making as there is a richer array of data to draw from when making decisions. In a scenario with additional data availability, we would expect the performance of CognitiveCharge to improve as the accuracy of the underlying analytics increases with greater information. Nevertheless, beyond the individual safety risk inherent to greater data sharing, there is an increased opportunity for malicious nodes to disseminate false information and disrupt the network. Currently, CognitiveCharge considers a single-hop ego network, and data is integrated only from nodes with which CognitiveCharge nodes come into direct contact. This increases CognitiveCharge nodes' robustness to false information; however, it is a limit on the scope of CognitiveCharge nodes' socio-spatio-temporal knowledge radius. Until a node visits a region directly and comes into contact with a node, its knowledge is restricted to its own area. Whilst this is expected for social-aware approaches, an extension to CognitiveCharge would be to explore multi-hop information dissemination as a strategy for wider and faster regional energy and threat awareness.

6.3 Applications

The VSG scenarios considered in our experimental analysis of CognitiveCharge focus on independent roaming EVs however our CognitiveCharge framework has broader applicability and deployment potential in a diverse range of scenarios. Managed EV fleets (and managed fleets in general) present an interesting area as CognitiveCharge can be combined with centralised and semi-centralised decision making to provide complementary means of managing energy.

Drones (UAVs) have seen rapidly increased adoption in recent years across a vast array of domains. UAVs present an interesting application area for CognitiveCharge due to the unique challenges drones face. At one end of the spectrum are small UAVs, such as the currently popular quadcopters, which are designed to be very lightweight and carry a small payload such as a digital camera or light parcel. At the other end of the spectrum are larger, more capable UAVs which are more commonly used in healthcare, industrial, and military domains. Regardless of size, use cases for UAVs most commonly fall under payload delivery or surveillance. Payload delivery ranges from rapid medication delivery to crop spreading whilst surveillance applications include areas such as traffic flow monitoring and stock inventory control [140].

The battery consumption rate of UAVs is strongly affected by the weight of the payload, which subsequently affects the usable flight time [141]. Much research has investigated the challenges of energy resource management for UAVs. In some respects, UAVs are similar to larger commercial EVs (such as lorries), where the payload represents a primary service and has the strongest impact on energy consumption. This is common for delivery services and is halting the industry drive towards EV adoption for large cargo transit (e.g. lorries) as the size of battery technologies eats into payload space. Unmanned underwater vehicles (UUVs) face similar but distinct challenges. Whilst EVs and UAVs support a wide variety of existing long range communications, UUVs are more limited by the environment [142]. CognitiveCharge has the potential to apply well to both UAV and UUV scenarios.

6.4 Simulations

In our experiments, we have modelled the characteristics of EVs using uniform values and represented their functions as linear. For instance, the rate at which EVs in our VSG scenario discharge energy through consumption for physical movement is uniform. Similarly, due to software limitations, the mobility of EVs is modelled on a 2D plane. These, and related necessary simplifications, have proven necessary for the experimental analysis of our proposed CognitiveCharge framework but are not the most precise representation of EVs in VSG scenarios. There are many additional factors that need to be taken into consideration for a more accurate representation of the VSG. These non-linearities are dependent upon a large number of factors, such as environmental conditions like temperature. For example, EV batteries are affected by real-world energy storage technologies that charge and discharge energy in a non-linear manner. It has been observed that many current battery technologies do not perform well in cold conditions [143]. Therefore, to accurately represent the behaviour of EVs in VSG scenarios, it is important to take into account this, and other factors which influence the performance of EVs in the real world. Similarly, the state of battery degradation and load dynamics are further examples of factors which may affect achieved performance. Our simulation approach takes reasonable steps to agnostically model and suitably represent EV characteristics on a macro scale and over a period of days without consideration to the specific micro conditions which would increase experimental complexity at the expense of scale. Nevertheless, this limits the experimental results to macro trends in our VSG scenarios as a lack of fidelity in simulations is well known to reduce accuracy [144].

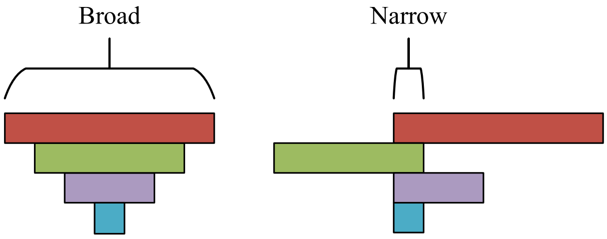

Figure 36 Scope and Wider Applicability

A particular challenge and limitation of this work lies in the scope of the VSG simulations conducted and broader applicability of results. Whilst the methodology employed is appropriate and made use of available data and suitable modelling to capture the cross-dimensional VSG dynamics as best possible for the context, the data is not fully complementary and necessitated being disjointly layered to derive suitable VSG scenarios. Figure 36 is a simple graphic which seeks to illustrate this. In an ideal case, where the built-up simulation layers are wholly complementary, the applicability of results extends much further than the core scenario. This is shown in the left image. Unfortunately, a lack of available cross-layer data for conducting experiments in our VSG contexts means that this work lies closer to the right image, with limitations in the wider applicability necessitating further development of novel scenarios to explore CognitiveCharge beyond this work.

Network simulations are extremely computationally intensive and current network simulation software is overwhelmingly restricted to single-core execution [145]. Performance improvement techniques widely employed in other disciplines for simulations include multi-threading and parallel processing paradigms using modern CPU and GPU architectures to drastically reduce real-world simulation time. Such techniques are extremely challenging to implement for network simulation software. Despite efforts over the years (e.g. [146]), only very recently have practical steps been realised in addressing these challenges [145]. The extremely time-consuming nature of conducting the experiments in this work has limited the ability to evaluate longer-term behaviours. A crucial avenue for future research in this area is to investigate the performance of CognitiveCharge and related protocols in large scale VSG scenarios where they can run for months of simulated time with many thousands of nodes. To capture the performance and offset limitations of various modelling strategies, we included several hybrid VSG simulation scenarios, which consisted of a variety of real-world, pseudo-realistic, and pseudorandom data. This allowed us to represent a diverse range of VSG scenarios. However, there are numerous opportunities available for further enhancing the realism of VSG simulations. One important area for future research is to develop new simulation software or improve existing ones to accurately depict the complexities and intricacies of VSG environments. By doing so, we can achieve a more comprehensive and realistic representation of VSG scenarios, leading to more accurate and insightful results. This is of particular benefit to emerging application domains, especially those where increased accuracy is more critical. This would be complemented by increasing availability of real-world, cross-layer datasets. For example, in the case of UAV payload delivery swarms, which have very short battery life and are highly affected by variables such as weather, temperature, altitude, and aerodynamics, the need for detailed and realistic environment modelling becomes even more vital. By incorporating such modelling, we can better understand and explore the real-world performance of these swarms across criteria such as communications and energy usage, ultimately leading to more informed decisions and improvements in their operations. Recent research has identified this as a particular need and simulation software is being actively developed which can capture the necessary fidelity [147], [148].

Through this work, the need for more advanced and comprehensive simulation tools became increasingly apparent. Whilst macro-scale trends and behaviours can be suitably modelled and explored, the lack of fidelity makes delving deeper into experimental results less meaningful. Simulation software plays a pivotal role in the development, testing, and optimisation of smart energy systems. In the coming years, we can anticipate the development of new, and the evolution of existing, simulation software that facilitates greater understanding of complex energy ecosystems. This includes not only the traditional elements of generation, distribution, and consumption but also the integration of renewable energy sources, energy storage solutions, and demand-side management. Enhanced simulation capabilities will enable researchers, engineers, and policymakers to model the intricate interplay of these elements with greater accuracy, allowing for more effective planning and decision-making. As SG paradigms such as VSGs become more prevalent, these tools will be indispensable in simulating and optimising their dynamic operations, ensuring grid resilience and adaptability. A vital area for future work is in realising performant simulation software with tooling to provide for realistic, large-scale modelling of VSGs. This necessitates in-depth modelling of the full scenario including communications, energy, network, mobility.

Chapter 7
Conclusion

7.1 Overview

This chapter is divided into three sections. This chapter begins by briefly summarising the work presented in this thesis. The research questions posed at the outset of this research are then revisited and the key findings of this thesis are highlighted alongside the contributions. Finally, concluding remarks and specific avenues for future work are given.

7.2 Thesis Summary

Chapter 1 introduced the work proving background on the smart grid (SG) and the vehicular smart grid (VSG) paradigms. Important assumptions of the VSG for this work were defined to set the scope of this thesis, and the energy denial-of-service attack (EDoS) threat model considered was detailed. The research questions central to this work were given, including the key sub-questions for investigation.

Chapter 2 offered a comprehensive review of literature related to the research focus of the thesis. The chapter began by identifying and detailing a set of criteria vital for approaches aiming to address self-organised energy resource awareness, P2P energy exchange, and adaptive resource-related security in vehicular smart grid environments. This was followed by a survey of existing works, with a primary focus on techniques for adaptive energy resource awareness, energy threat context awareness, and reputation-based security in opportunistic networks.

Chapter 3 introduced our CognitiveCharge proposal, detailing how it detected and responded in real-time to dynamic in-network conditions spanning mobility, resources, services, and security. An architectural overview was provided that highlighted the key challenges and conflicting trade-offs in this intricate design space. The analytical and functional models of our CognitiveCharge framework were also given. Each of the components of CognitiveCharge were described and algorithms provided for the core components: the cross-layer predictive analytics, collaborative peer testing mechanism, and real-time utility-driven decision-making processes.

Chapter 4 presented the experimental methodology used for a rigorous analysis and evaluation of the CognitiveCharge proposal. This began with a description of the simulation environments and an explanation of the hybrid real-world and pseudorealistic data traces representing each dimension of the problem scenario. The VSG scenarios used in this work were described, and a comparative analysis of these scenarios was provided. Subsequently, the modelling of the energy depletion attack in these scenarios was detailed.

Chapter 5 contained an evaluation of the implemented prototype of the CognitiveCharge framework using the experimental methodology detailed in Chapter 4. This chapter used a variety of real-world, pseudo-realistic, and pseudo-random VSG scenarios to understand the performance of the CognitiveCharge proposal under different conditions. Extensive evaluations occurred in vehicular smart grid scenarios with fully trusted nodes, mutually untrusted nodes, and nodes under active energy depletion attacks by injected malicious entities. Under each condition, performance characteristics for CognitiveCharge were compared with baseline conditions and benchmarked against other approaches.

Chapter 6 examined the wider context of our CognitiveCharge framework, discussing the broader picture of the proposal presented in this thesis. By taking into consideration the criteria set outlined in Chapter 2, Chapter 6 looked at the feasibility and challenges of real-world deployments of CognitiveCharge in domains with varying constraints and dynamics. This included deployments across alternative topologies with different degrees of centralisation and heterogeneity. In particular, Chapter 6 was an opportunity to explore alternative application domains for our proposed CognitiveCharge framework as well as scenarios where the criteria set given in Chapter 2 applied less stringently. Chapter 6 also looked at specific aspects of this work in greater depth and gave consideration to avenues for future research.

7.3 Research Questions

In Chapter 1, the overall research question of this thesis was introduced and motivated. This was subsequently broken down into three key sub-questions for investigation in this work. The overall research question was as given as follows:

With fully localised communication and decision-making, is it possible to increase the utility of nodes and limit energy losses in the presence of an energy depletion attack conducted by malicious nodes in heterogeneous VSGs?

The sub-questions are also repeated here, numbered for easier reference. These will be subsequently referred to as RQ1, RQ2, and RQ3. Each of the given research questions has been addressed in specific areas of this work which are outlined below.

1. *How does fully localised communication and decision-making affect the utility of nodes in heterogeneous VSGs?*
2. *What are the energy losses incurred by an EDoS attack and to what extent can this be detected and reduced?*
3. *What are the associated trade-offs between peer trust and energy availability?*

This thesis has been carefully structured to demonstrate how each of these research questions have been considered and addressed. Regarding the overall research question, the key criteria were first identified in Chapter 2 and used to conduct a literature review of relevant works. Chapter 3 proposed our novel CognitiveCharge framework which meets these criteria and comprises a suite of novel predictive analytics and collaborative peer-testing scheme combined by a real-time utility driven decision making. In Chapter 5 we explored the performance advantages and disadvantages of CognitiveCharge in untrusted VSGs and compared CognitiveCharge against a range of similar localised decision-making approaches. The overall research question and the context of the research was considered more broadly in Chapter 6.

RQ1 focuses specifically on exploring and understanding the feasibility and impacts of localised decision making on EVs in VSG environments. In the first phase of our experiments, we investigated the performance of localised decision making on criteria relating to the energy-resource dynamics and behaviour of EVs in a range of VSG scenarios without presence of malicious nodes. These directly link with our energy-resource aware CognitiveCharge analytics and the relevant criteria identified in Chapter 2.

RQ2 looks specifically at the EDoS attack central to this work. Robustness to attack is a core design consideration of our CognitiveCharge framework detailed in Chapter 3, and we further consider a peer testing and reputation exchange mechanism which seeks to directly detect and mitigate the impact of EDoS attacks on EVs in the VSG. In Chapter 5, the performance of these approaches is explored across measures, including responsiveness and accuracy.

RQ3 brings together RQ1 and RQ2 to consider the trade-offs between adaptive, localised energy management and security. Our CognitiveCharge framework is designed to adaptively balance these trade-offs so that CognitiveCharge nodes can continue to exchange energy, even in the presence of EDoS attack. Chapter 5 explores the trade-off between energy and security in multiple VSG scenarios.

7.4 Findings and Contributions

This thesis has made modest but important contributions to the field of computer networks and smart energy. In the context of VSGs, this work has proposed a novel framework extended and explored novel adaptation and application of established approaches to the VSG context. Through simulation-based experimental analysis, we found that the proposals made in the work have benefit to the energy-resource utility of nodes in VSGs. We further find that our CognitiveCharge proposal is a viable framework for detecting and mitigating EDoS attacks in VSGs. As discussed in Chapter 6, consideration has been given to the broader policy and industry implications of CognitiveCharge. This included a deep examination of the potential avenues for extension that may arise in various domains. Limitations of CognitiveCharge and this thesis in general have been identified throughout and attention has been given to directions for future work to overcome these and offer advancements to CognitiveCharge. The key contributions of this work are as follows:

- A suite of novel, cross-layer, predictive energy and threat context-aware analytics for capturing and interpreting the VSG environment from the perspective of a VSG node.
- A novel framework for analytics exchange, peer data integration, and real-time energy-resource utility-driven decision-making in untrusted VSGs.
- A proactive, collaborative peer testing mechanism for energy-resource behaviour evaluation in untrusted VSGs for detection and mitigation of EDoS attack.

The core contributions outlined have benefit beyond the VSG context and can be explored in other domains. The additional contributions of this work lie in the exploration of CognitiveCharge, and the related strategies investigated in the context of VSG scenarios comprising EVs and CSs. Further contributions arise from the reusable software generated and data collected for this work. The experimentation and analysis conducted offer valuable insights into the management of energy resources in untrusted opportunistic VSGs, both with and without malicious EDoS attackers.

7.5 Future Work

Areas for future work and the relevance of this work to other research is discussed broadly in Chapter 6. Here we highlight several specific avenues for research to follow on from this thesis.

There is a significant, pressing need for extension to existing, or entirely new, simulation software which is capable of modelling the complexities of VSG environments and has sufficient performance to conduct these experiments at scale. This will be increasingly important in coming years as rapidly increasing adoption of EVs, and related technologies such as UAVs, become increasingly widespread, get deployed in new contexts, and the challenges identified in the literature, such as grid stability and risk of attack, become increasingly prevalent.

Unmanned vehicles (autonomous EVs, UAVs, UUVs) present interesting future domains for exploring CognitiveCharge and related works. Autonomous EVs are slowly seeing test-based deployments in various locations around the world. UAVs have seen rapid adoption worldwide and increasing usage across a vast array of domains. These environments are often highly challenging and resource constrained, necessitate real-time decision making. Exploring fully-localised security and energy aware approaches in these environments would benefit the networks themselves as well as feeding back direct improvements to CognitiveCharge and other approaches for the benefit of other application domains.

Beyond the SG environment, there are several interesting opportunities for future work to explore feeding the concepts presented in this work into closely related areas where networking principles can be apply both directly and indirectly. This work applied opportunistic networking (OppNet) theories to VSGs and the P2P routing of energy amongst mutually untrusted nodes. In a previous work, we explored this concept but with physical resources in the aftermath of disaster; for instance, so that individuals and groups can share and acquire distributed resources such as medication and equipment [149]. Energy is a special case of resource as whilst it can be acquired and exchanged, it can also be generated. In this work we particularly focused on independent EVs collaboratively maintaining service provision in the presence of an active EDoS attack. In past works we have considered the protracted management of resources in disaster scenarios [149]. A key area for future work would be to consider in more depth the transitional and interim periods between disaster stages. Research with this level of fidelity is largely underexplored in the context of energy exchange amongst EVs.

The attack central to this work is an EDoS attack conducted by actively attacking malicious nodes. An area of further investigation is to explore the performance of CognitiveCharge, and related approaches, under similar attacks and in more challenging scenarios. For EDoS attacks in particular, it is important to consider in more depth the impacts of combinations of passive and attacks and cases where nodes seek to mask their behaviour to avoid being detected, the effectiveness of strategies to reduce the impacts of such attacks, and finally to explore robustness of approaches in the cases where attackers look to directly exploit threat detection and collaborative trust strategies.

References

[1] J. Zhou, L. He, C. Li, Y. Cao, X. Liu, and Y. Geng, 'What's the difference between traditional power grid and smart grid? — From dispatching perspective', in *2013 IEEE PES Asia-Pacific Power and Energy Engineering Conference (APPEEC)*, Dec. 2013, pp. 1–6.

[2] S. Skippon and M. Garwood, 'Responses to battery electric vehicles: UK consumer attitudes and attributions of symbolic meaning following direct experience to reduce psychological distance', *Transp. Res. Part Transp. Environ.*, vol. 16, no. 7, pp. 525–531, Oct. 2011,

[3] L. Bunce, M. Harris, and M. Burgess, 'Charge up then charge out? Drivers' perceptions and experiences of electric vehicles in the UK', *Transp. Res. Part Policy Pract.*, vol. 59, pp. 278–287, Jan. 2014,

[4] S. W. Hadley and A. A. Tsvetkova, 'Potential Impacts of Plug-in Hybrid Electric Vehicles on Regional Power Generation', *Electr. J.*, vol. 22, no. 10, pp. 56–68, Dec. 2009,

[5] M. Ashfaq, O. Butt, J. Selvaraj, and N. Rahim, 'Assessment of electric vehicle charging infrastructure and its impact on the electric grid: A review', *Int. J. Green Energy*, vol. 18, no. 7, pp. 657–686, May 2021,

[6] P. K. Joseph, E. Devaraj, and A. Gopal, 'Overview of wireless charging and vehicle-to-grid integration of electric vehicles using renewable

energy for sustainable transportation', *IET Power Electron.*, vol. 12, no. 4, pp. 627–638, Dec. 2018,

[7] M. A. Ponce-Jara, E. Ruiz, R. Gil, E. Sancristóbal, C. Pérez-Molina, and M. Castro, 'Smart Grid: Assessment of the past and present in developed and developing countries', *Energy Strategy Rev.*, vol. 18, pp. 38–52, Dec. 2017,

[8] C. Wueest, 'Targeted Attacks Against the Energy Sector', Symantec, Jan. 2014.

[9] H. Zhang, B. Liu, and H. Wu, 'Smart Grid Cyber-Physical Attack and Defense: A Review', *IEEE Access*, vol. 9, pp. 29641–29659, 2021,

[10] D. B. Richardson, 'Electric vehicles and the electric grid: A review of modeling approaches, Impacts, and renewable energy integration', *Renew. Sustain. Energy Rev.*, vol. 19, pp. 247–254, Mar. 2013,

[11] P. Antmann, 'Reducing Technical and Non-Technical Losses in the Power Sector', World Bank Group, 2009.

[12] S. S. S. R. Depuru, L. Wang, and V. Devabhaktuni, 'Electricity theft: Overview, issues, prevention and a smart meter based approach to control theft', *Energy Policy*, vol. 39, no. 2, pp. 1007–1015, Feb. 2011,

[13] N. Kshetri and J. Voas, 'Hacking Power Grids: A Current Problem', *Computer*, vol. 50, no. 12, pp. 91–95, Dec. 2017,

[14] D. E. Whitehead, K. Owens, D. Gammel, and J. Smith, 'Ukraine cyber-induced power outage: Analysis and practical mitigation strategies', in *2017 70th Annual Conference for Protective Relay Engineers (CPRE)*, Apr. 2017, pp. 1–8.

[15] E. U. Soykan, M. Bagriyanik, and G. Soykan, ‘Disrupting the power grid via EV charging: The impact of the SMS Phishing attacks’, *Sustain. Energy Grids Netw.*, vol. 26, p. 100477, Jun. 2021,

[16] S. Saha, A. Lukyanenko, and A. Ylä-Jääski, ‘Cooperative caching through routing control in information-centric networks’, in *2013 Proceedings IEEE INFOCOM*, Apr. 2013, pp. 100–104.

[17] D. Bertsimas, V. F. Farias, and N. Trichakis, ‘The Price of Fairness’, *Oper. Res.*, vol. 59, no. 1, pp. 17–31, Feb. 2011,

[18] G. N. Sorebo, M. C. Echols, and M. C. Echols, *Smart Grid Security : An End-to-End View of Security in the New Electrical Grid*. CRC Press, 2011.

[19] X. Fang, S. Misra, G. Xue, and D. Yang, ‘Smart Grid — The New and Improved Power Grid: A Survey’, *IEEE Commun. Surv. Tutor.*, vol. 14, no. 4, pp. 944–980, Fourth 2012,

[20] S. Massoud Amin and B. F. Wollenberg, ‘Toward a smart grid: power delivery for the 21st century’, *IEEE Power Energy Mag.*, vol. 3, no. 5, pp. 34–41, Sep. 2005,

[21] T. Li, H. Zhao, S. Wang, C. Yang, and B. Huang, ‘Attack and Defense Strategy of Distribution Network Cyber-Physical System Considering EV Source-Charge Bidirectionality’, *Electronics*, vol. 10, no. 23, Art. no. 23, Jan. 2021,

[22] G. Zyba, G. M. Voelker, S. Ioannidis, and C. Diot, ‘Dissemination in opportunistic mobile ad-hoc networks: The power of the crowd’, in *2011 Proceedings IEEE INFOCOM*, Apr. 2011, pp. 1179–1187.

[23] A. M. Vegni and V. Loscrí, ‘A Survey on Vehicular Social Networks’, *IEEE Commun. Surv. Tutor.*, vol. 17, no. 4, pp. 2397–2419, Fourthquarter 2015,

[24] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, ‘Pocket switched networks and human mobility in conference environments’, in *Proceedings of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking*, in WDTN ’05. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, Aug. 2005, pp. 244–251.

[25] U. Brandes and T. Erlebach, Eds., *Network Analysis: Methodological Foundations*. in Theoretical Computer Science and General Issues. Berlin Heidelberg: Springer-Verlag, 2005.

[26] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, ‘Impact of Human Mobility on Opportunistic Forwarding Algorithms’, *IEEE Trans. Mob. Comput.*, vol. 6, no. 6, pp. 606–620, Jun. 2007,

[27] A. Vahdat and D. Becker, ‘Epidemic Routing for Partially-Connected Ad Hoc Networks’, 2000.

[28] P. V. Marsden, ‘Egocentric and sociocentric measures of network centrality’, *Soc. Netw.*, vol. 24, no. 4, pp. 407–422, Oct. 2002,

[29] V. Erramilli and M. Crovella, ‘Forwarding in opportunistic networks with resource constraints’, in *Proceedings of the third ACM workshop on Challenged networks*, in CHANTS ’08. New York, NY, USA: Association for Computing Machinery, Sep. 2008, pp. 41–48.

[30] H. Dubois-Ferriere, M. Grossglauser, and M. Vetterli, ‘Age matters: efficient route discovery in mobile ad hoc networks using encounter ages’, in *Proceedings of the 4th ACM international symposium on Mobile*

ad hoc networking & computing, in MobiHoc '03. New York, NY, USA: Association for Computing Machinery, Jun. 2003, pp. 257–266.

[31] A. Lindgren, A. Doria, and O. Schelén, 'Probabilistic routing in intermittently connected networks', *ACM SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 7, no. 3, pp. 19–20, Jul. 2003,

[32] A. Lindgren, A. Doria, and O. Schelén, 'Probabilistic Routing in Intermittently Connected Networks', in *Service Assurance with Partial and Intermittent Resources*, P. Dini, P. Lorenz, and J. N. de Souza, Eds., in Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2004, pp. 239–254.

[33] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine, 'MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks', in *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, Apr. 2006, pp. 1–11.

[34] V. Erramilli, M. Crovella, A. Chaintreau, and C. Diot, 'Delegation forwarding', in *Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, in MobiHoc '08. New York, NY, USA: Association for Computing Machinery, May 2008, pp. 251–260.

[35] A. N. Shiryaev, *Optimal Stopping Rules*. in Stochastic Modelling and Applied Probability. Berlin Heidelberg: Springer-Verlag, 2008.

[36] P. Hui and J. Crowcroft, 'How Small Labels Create Big Improvements', in *Fifth Annual IEEE International Conference on Pervasive Computing and Communications Workshops (PerComW'07)*, Mar. 2007, pp. 65–70.

[37] P. Hui, J. Crowcroft, and E. Yoneki, ‘Bubble rap: social-based forwarding in delay tolerant networks’, in *Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, in MobiHoc ’08. New York, NY, USA: Association for Computing Machinery, May 2008, pp. 241–250.

[38] E. M. Daly and M. Haahr, ‘Social Network Analysis for Information Flow in Disconnected Delay-Tolerant MANETs’, *IEEE Trans. Mob. Comput.*, vol. 8, no. 5, pp. 606–621, May 2009,

[39] T.-K. Huang, C.-K. Lee, and L.-J. Chen, ‘PRoPHET+: An Adaptive PRoPHET-Based Routing Protocol for Opportunistic Network’, in *2010 24th IEEE International Conference on Advanced Information Networking and Applications*, Apr. 2010, pp. 112–119.

[40] M. Radenkovic and A. Grundy, ‘Congestion aware forwarding in delay tolerant and social opportunistic networks’, in *2011 Eighth International Conference on Wireless On-Demand Network Systems and Services*, Jan. 2011, pp. 60–67.

[41] M. Radenkovic, V. S. H. Huynh, and P. Manzoni, ‘Adaptive Real-Time Predictive Collaborative Content Discovery and Retrieval in Mobile Disconnection Prone Networks’, *IEEE Access*, vol. 6, pp. 32188–32206, 2018,

[42] V. S. H. Huynh and M. Radenkovic, ‘A novel cross-layer framework for large scale emergency communications’, in *2017 13th International Wireless Communications and Mobile Computing Conference (IWCMC)*, Jun. 2017, pp. 2152–2157.

[43] S. D. Han and Y. W. Chung, 'An Improved PRoPHET Routing Protocol in Delay Tolerant Network', The Scientific World Journal.

[44] K. Wei, X. Liang, and K. Xu, 'A Survey of Social-Aware Routing Protocols in Delay Tolerant Networks: Applications, Taxonomy and Design-Related Issues', *IEEE Commun. Surv. Tutor.*, vol. 16, no. 1, pp. 556–578, 2014,

[45] C. Binding *et al.*, 'Electric vehicle fleet integration in the danish EDISON project - A virtual power plant on the island of Bornholm', in *IEEE PES General Meeting*, Jul. 2010, pp. 1–8.

[46] S. Zou, Z. Ma, X. Liu, and I. Hiskens, 'An Efficient Game for Coordinating Electric Vehicle Charging', *IEEE Trans. Autom. Control*, vol. 62, no. 5, pp. 2374–2389, May 2017,

[47] Z. Ma, D. S. Callaway, and I. A. Hiskens, 'Decentralized Charging Control of Large Populations of Plug-in Electric Vehicles', *IEEE Trans. Control Syst. Technol.*, vol. 21, no. 1, pp. 67–78, Jan. 2013,

[48] L. Gan, U. Topcu, and S. H. Low, 'Optimal decentralized protocol for electric vehicle charging', *IEEE Trans. Power Syst.*, vol. 28, no. 2, pp. 940–951, May 2013,

[49] A. Ghavami, K. Kar, and A. Gupta, 'Decentralized Charging of Plug-in Electric Vehicles With Distribution Feeder Overload Control', *IEEE Trans. Autom. Control*, vol. 61, no. 11, pp. 3527–3532, Nov. 2016,

[50] L. A. Maglaras, F. V. Topalis, and A. L. Maglaras, 'Cooperative approaches for dymanic wireless charging of Electric Vehicles in a smart city', in *2014 IEEE International Energy Conference (ENERGYCON)*, May 2014, pp. 1365–1369.

[51] M. Wang, M. Ismail, R. Zhang, X. S. Shen, E. Serpedin, and K. Qaraqe, ‘A semi-distributed V2V fast charging strategy based on price control’, in *2014 IEEE Global Communications Conference*, Dec. 2014, pp. 4550–4555.

[52] P. Dutta, ‘Charge sharing model using inductive power transfer to increase feasibility of electric vehicle taxi fleets’, in *2014 IEEE PES General Meeting | Conference Exposition*, Jul. 2014, pp. 1–4.

[53] E. Bulut and M. C. Kisacikoglu, ‘Mitigating Range Anxiety via Vehicle-to-Vehicle Social Charging System’, in *2017 IEEE 85th Vehicular Technology Conference (VTC Spring)*, Jun. 2017, pp. 1–5.

[54] M. R. Sarker, H. Pandžić, and M. A. Ortega-Vazquez, ‘Optimal Operation and Services Scheduling for an Electric Vehicle Battery Swapping Station’, *IEEE Trans. Power Syst.*, vol. 30, no. 2, pp. 901–910, Mar. 2015,

[55] R. Alvaro-Hermana, J. Fraile-Ardanuy, P. J. Zufiria, L. Knapen, and D. Janssens, ‘Peer to Peer Energy Trading with Electric Vehicles’, *IEEE Intell. Transp. Syst. Mag.*, vol. 8, no. 3, pp. 33–44, Fall 2016,

[56] J. Wu, Y. Zhu, L. Liu, B. Yu, and J. Pan, ‘Energy-Efficient Routing in Multi-Community DTN with Social Selfishness Considerations’, in *2016 IEEE Global Communications Conference (GLOBECOM)*, Dec. 2016, pp. 1–7.

[57] N. Banerjee, M. D. Corner, and B. N. Levine, ‘An Energy-Efficient Architecture for DTN Throwboxes’, in *IEEE INFOCOM 2007 - 26th IEEE International Conference on Computer Communications*, May 2007, pp. 776–784.

[58] E. Spaho, K. Dhoska, K. Bylykbashi, L. Barolli, V. Kolici, and M. Takizawa, ‘Performance Evaluation of Energy Consumption for Different DTN Routing Protocols’, in *Advances in Network-Based Information Systems*, L. Barolli, N. Kryvinska, T. Enokido, and M. Takizawa, Eds., Cham: Springer International Publishing, 2019, pp. 122–131.

[59] M. A. Rahmadhani, L. V. Yovita, and R. Mayasari, ‘Energy Consumption and Packet Loss Analysis of LEACH Routing Protocol on WSN Over DTN’, in *2018 4th International Conference on Wireless and Telematics (ICWT)*, Jul. 2018, pp. 1–5.

[60] D. Vardalis and V. Tsaoussidis, ‘Exploiting the potential of DTN for energy-efficient internetworking’, *J. Syst. Softw.*, vol. 90, pp. 91–103, Apr. 2014,

[61] M. Marchese and F. Patrone, ‘Energy-aware Routing Algorithm for DTN-Nanosatellite Networks’, in *2018 IEEE Global Communications Conference (GLOBECOM)*, Dec. 2018, pp. 206–212.

[62] Q. Yan, B. Zhang, and M. Kezunovic, ‘Optimized Operational Cost Reduction for an EV Charging Station Integrated With Battery Energy Storage and PV Generation’, *IEEE Trans. Smart Grid*, vol. 10, no. 2, pp. 2096–2106, Mar. 2019,

[63] W. Jiang and Y. Zhen, ‘A Real-Time EV Charging Scheduling for Parking Lots With PV System and Energy Store System’, *IEEE Access*, vol. 7, pp. 86184–86193, 2019,

[64] ‘Optimized Charge Scheduling of Electric Buses in a City Bus Network | IEEE Conference Publication | IEEE Xplore’.

[65] J. Liu, G. Lin, S. Huang, Y. Zhou, C. Rehtanz, and Y. Li, ‘Collaborative EV Routing and Charging Scheduling With Power Distribution and Traffic Networks Interaction’, *IEEE Trans. Power Syst.*, vol. 37, no. 5, pp. 3923–3936, Sep. 2022,

[66] U. Baroudi, ‘Robot-Assisted Maintenance of Wireless Sensor Networks Using Wireless Energy Transfer’, *IEEE Sens. J.*, vol. 17, no. 14, pp. 4661–4671, Jul. 2017,

[67] S. Wang *et al.*, ‘Robotic Wireless Energy Transfer in Dynamic Environments: System Design and Experimental Validation’, *IEEE Commun. Mag.*, vol. 60, no. 3, pp. 40–46, Mar. 2022,

[68] RAC, ‘EV Boost Mobile Charging System’.

[69] S. C. Nelson, M. Bakht, and R. Kravets, ‘Encounter-Based Routing in DTNs’, in *IEEE INFOCOM 2009*, Apr. 2009, pp. 846–854.

[70] M. Radenkovic and I. Vaghi, ‘Adaptive User Anonymity for Mobile Opportunistic Networks’, in *Proceedings of the Seventh ACM International Workshop on Challenged Networks*, in CHANTS ’12. New York, NY, USA: ACM, 2012, pp. 79–82.

[71] M. Radenkovic, A. Benslimane, and D. McAuley, ‘Reputation Aware Obfuscation for Mobile Opportunistic Networks’, *IEEE Trans. Parallel Distrib. Syst.*, vol. 26, no. 1, pp. 230–240, Jan. 2015,

[72] M. Radenkovic, ‘Cognitive Privacy for Personal Clouds’, Mobile Information Systems.

[73] Y. Gao, J. Tao, Y. Xu, Z. Wang, W. Sun, and G. Cheng, ‘CEBD: Contact-Evidence-Driven Blackhole Detection Based on Machine Learning in

OppNets', *IEEE Trans. Comput. Soc. Syst.*, vol. 8, no. 6, pp. 1344–1356, Dec. 2021,

[74] A. Chhabra, V. Vashishth, and D. K. Sharma, 'A fuzzy logic and game theory based adaptive approach for securing opportunistic networks against black hole attacks', *Int. J. Commun. Syst.*, vol. 31, no. 4, p. e3487, 2018,

[75] R. T. Merlin and R. Ravi, 'Novel Trust Based Energy Aware Routing Mechanism for Mitigation of Black Hole Attacks in MANET', *Wirel. Pers. Commun.*, vol. 104, no. 4, pp. 1599–1636, Feb. 2019,

[76] V. Kumar and R. Kumar, 'A Cooperative Black Hole Node Detection and Mitigation Approach for MANETs', in *Innovative Security Solutions for Information Technology and Communications*, I. Bica, D. Naccache, and E. Simion, Eds., in Lecture Notes in Computer Science. Cham: Springer International Publishing, 2015, pp. 171–183.

[77] K. J. Sarma, R. Sharma, and R. Das, 'A survey of Black hole attack detection in Manet', in *2014 International Conference on Issues and Challenges in Intelligent Computing Techniques (ICICT)*, Feb. 2014, pp. 202–205.

[78] S. Gurung and S. Chauhan, 'A survey of black-hole attack mitigation techniques in MANET: merits, drawbacks, and suitability', *Wirel. Netw.*, vol. 26, no. 3, pp. 1981–2011, Apr. 2020,

[79] F.-H. Tseng, L.-D. Chou, and H.-C. Chao, 'A survey of black hole attacks in wireless mobile ad hoc networks', *Hum.-Centric Comput. Inf. Sci.*, vol. 1, no. 1, Art. no. 1, Dec. 2011,

[80] C. Molina-Jimenez, E. Solaiman, I. Sfyrakis, I. Ng, and J. Crowcroft, ‘On and Off-Blockchain Enforcement of Smart Contracts’, in *Euro-Par 2018: Parallel Processing Workshops*, G. Mencagli, D. B. Heras, V. Cardellini, E. Casalicchio, E. Jeannot, F. Wolf, A. Salis, C. Schifanella, R. R. Manumachu, L. Ricci, M. Beccuti, L. Antonelli, J. D. Garcia Sanchez, and S. L. Scott, Eds., in Lecture Notes in Computer Science. Springer International Publishing, 2019, pp. 342–354.

[81] J. Poon and T. Dryja, *The bitcoin lightning network: Scalable off-chain instant payments*. 2016.

[82] M. U. Hassan, M. H. Rehmani, and J. Chen, ‘Privacy preservation in blockchain based IoT systems: Integration issues, prospects, challenges, and future research directions’, *Future Gener. Comput. Syst.*, Mar. 2019,

[83] E. Mengelkamp, J. Gärttner, K. Rock, S. Kessler, L. Orsini, and C. Weinhardt, ‘Designing microgrid energy markets: A case study: The Brooklyn Microgrid’, *Appl. Energy*, vol. 210, pp. 870–880, Jan. 2018,

[84] M. Mihaylov, S. Jurado, K. V. Moffaert, N. Avellana, and A. Nowé, ‘NRG-X-Change - A Novel Mechanism for Trading of Renewable Energy in Smart Grids’, in *SMARTGREENS*, 2014.

[85] M. Mihaylov, S. Jurado, N. Avellana, K. V. Moffaert, I. M. de Abril, and A. Nowé, ‘NRGcoin: Virtual currency for trading of renewable energy in smart grids’, in *11th International Conference on the European Energy Market (EEM14)*, May 2014, pp. 1–6.

[86] M. Mihaylov, I. Razo-Zapata, R. Rădulescu, S. Jurado, N. Avellana, and A. Nowé, ‘Smart Grid Demonstration Platform for Renewable Energy Exchange’, in *Advances in Practical Applications of Scalable Multi-*

agent Systems. The PAAMS Collection, Y. Demazeau, T. Ito, J. Bajo, and M. J. Escalona, Eds., in Lecture Notes in Computer Science. Springer International Publishing, 2016, pp. 277–280.

[87] A. Laszka, A. Dubey, M. Walker, and D. Schmidt, ‘Providing Privacy, Safety, and Security in IoT-based Transactive Energy Systems Using Distributed Ledgers’, in *Proceedings of the Seventh International Conference on the Internet of Things*, in IoT ’17. New York, NY, USA: ACM, 2017, p. 13:1-13:8.

[88] N. Z. Aitzhan and D. Svetinovic, ‘Security and Privacy in Decentralized Energy Trading Through Multi-Signatures, Blockchain and Anonymous Messaging Streams’, *IEEE Trans. Dependable Secure Comput.*, vol. 15, no. 5, pp. 840–852, Sep. 2018,

[89] Z. Su, Y. Wang, Q. Xu, M. Fei, Y. Tian, and N. Zhang, ‘A Secure Charging Scheme for Electric Vehicles with Smart Communities in Energy Blockchain’, *IEEE Internet Things J.*, pp. 1–1, 2019,

[90] J. Kang, R. Yu, X. Huang, S. Maharjan, Y. Zhang, and E. Hossain, ‘Enabling Localized Peer-to-Peer Electricity Trading Among Plug-in Hybrid Electric Vehicles Using Consortium Blockchains’, *IEEE Trans. Ind. Inform.*, vol. 13, no. 6, pp. 3154–3164, Dec. 2017,

[91] C. B. Avoussoukpo, T. B. Ogunseyi, and M. Tchenagnon, ‘Securing and Facilitating Communication Within Opportunistic Networks: A Holistic Survey’, *IEEE Access*, vol. 9, pp. 55009–55035, 2021,

[92] R. Sachdeva and A. Dev, ‘Review of opportunistic network: Assessing past, present, and future’, *Int. J. Commun. Syst.*, vol. 34, no. 11, p. e4860, 2021,

[93] C. Chakrabarti, ‘A trust based scheme to detect selfish nodes using delay tolerant network’, *Glob. J. Res. Anal.*, vol. 8, no. 6, Jul. 2019.

[94] R. Hussain, J. Lee, and S. Zeadally, ‘Trust in VANET: A Survey of Current Solutions and Future Research Opportunities’, *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 5, pp. 2553–2571, May 2021,

[95] D. Shehada, A. Gawanmeh, C. Y. Yeun, and M. Jamal Zemerly, ‘Fog-based distributed trust and reputation management system for internet of things’, *J. King Saud Univ. - Comput. Inf. Sci.*, vol. 34, no. 10, Part A, pp. 8637–8646, Nov. 2022,

[96] D. Shehada, C. Y. Yeun, M. Jamal Zemerly, M. Al-Qutayri, Y. Al-Hammadi, and J. Hu, ‘A new adaptive trust and reputation model for Mobile Agent Systems’, *J. Netw. Comput. Appl.*, vol. 124, pp. 33–43, Dec. 2018,

[97] H. Yu, Z. Shen, C. Miao, C. Leung, and D. Niyato, ‘A Survey of Trust and Reputation Management Systems in Wireless Communications’, *Proc. IEEE*, vol. 98, no. 10, pp. 1755–1772, Oct. 2010,

[98] H. Xia, S. Zhang, B. Li, L. Li, and X. Cheng, ‘Towards a Novel Trust-Based Multicast Routing for VANETs’, *Secur. Commun. Netw.*, vol. 2018, p. e7608198, Oct. 2018,

[99] P. Asuquo, H. Cruickshank, C. P. A. Ogah, A. Lei, and Z. Sun, ‘A Distributed Trust Management Scheme for Data Forwarding in Satellite DTN Emergency Communications’, *IEEE J. Sel. Areas Commun.*, vol. 36, no. 2, pp. 246–256, Feb. 2018,

[100] P. Asuquo, H. Cruickshank, C. P. Anyigor Ogah, A. Lei, and K. Olutomilayo, ‘A Mobility-Aware Trust Management Scheme for

Emergency Communication Networks Using DTN', in *Wireless and Satellite Systems*, I. Otung, P. Pillai, G. Eleftherakis, and G. Giambene, Eds., Cham: Springer International Publishing, 2017, pp. 130–141.

[101] X. Zhang, X. Wang, A. Liu, Q. Zhang, and C. Tang, 'Reputation-based scheme for delay tolerant networks', in *Proceedings of 2011 International Conference on Computer Science and Network Technology*, Dec. 2011, pp. 974–978.

[102] P. Asuquo, H. Cruickshank, C. P. Anyigor Ogah, A. Lei, and Z. Sun, 'A collaborative trust management scheme for emergency communication using delay tolerant networks', in *2016 8th Advanced Satellite Multimedia Systems Conference and the 14th Signal Processing for Space Communications Workshop (ASMS/SPSC)*, Sep. 2016, pp. 1–6.

[103] A. Grundy and M. Radenkovic, 'Promoting congestion control in opportunistic networks', in *2010 IEEE 6th International Conference on Wireless and Mobile Computing, Networking and Communications*, Oct. 2010, pp. 324–330.

[104] M. Radenkovic and V. S. Ha Huynh, 'Energy-Aware Opportunistic Charging and Energy Distribution for Sustainable Vehicular Edge and Fog Networks', in *2020 Fifth International Conference on Fog and Mobile Edge Computing (FMEC)*, Apr. 2020, pp. 5–12.

[105] J. Tang, M. Musolesi, C. Mascolo, and V. Latora, 'Characterising Temporal Distance and Reachability in Mobile and Online Social Networks', *SIGCOMM Comput Commun Rev*, vol. 40, no. 1, pp. 118–124, Jan. 2010,

[106] Y. Wang, Y. Yuan, Y. Ma, and G. Wang, ‘Time-Dependent Graphs: Definitions, Applications, and Algorithms’, *Data Sci. Eng.*, vol. 4, no. 4, pp. 352–366, Dec. 2019,

[107] A. Grundy, ‘Congestion control framework for delay-tolerant communications’, Ph.D., University of Nottingham, 2012.

[108] M. Radenkovic and A. Walker, ‘CognitiveCharge: Disconnection Tolerant Adaptive Collaborative and Predictive Vehicular Charging’, in *Proceedings of the 4th ACM MobiHoc Workshop on Experiences with the Design and Implementation of Smart Objects*, in SMARTOBJECTS ’18. New York, NY, USA: ACM, 2018, p. 2:1-2:9.

[109] E. B. Wilson, ‘Probable Inference, the Law of Succession, and Statistical Inference’, *J. Am. Stat. Assoc.*, vol. 22, no. 158, pp. 209–212, Jun. 1927,

[110] A. Agresti and B. A. Coull, ‘Approximate is Better than “Exact” for Interval Estimation of Binomial Proportions’, *Am. Stat.*, vol. 52, no. 2, pp. 119–126, May 1998,

[111] E. Weingartner, H. vom Lehn, and K. Wehrle, ‘A Performance Comparison of Recent Network Simulators’, in *2009 IEEE International Conference on Communications*, Jun. 2009, pp. 1–5.

[112] I. Minakov, R. Passerone, A. Rizzardi, and S. Sicari, ‘A Comparative Study of Recent Wireless Sensor Network Simulators’, *ACM Trans Sen Netw*, vol. 12, no. 3, p. 20:1-20:39, Jul. 2016,

[113] R. Sharma, V. Vashisht, and U. Singh, ‘Modelling and simulation frameworks for wireless sensor networks: a comparative study’, *IET Wirel. Sens. Syst.*, vol. 10, no. 5, pp. 181–197, Oct. 2020,

[114] M. Radenkovic, J. Crowcroft, and M. H. Rehmani, 'Towards Low Cost Prototyping of Mobile Opportunistic Disconnection Tolerant Networks and Systems', *IEEE Access*, vol. 4, pp. 5309–5321, 2016,

[115] X. Zeng, R. Bagrodia, and M. Gerla, 'GloMoSim: a library for parallel simulation of large-scale wireless networks', in *Proceedings. Twelfth Workshop on Parallel and Distributed Simulation PADS '98 (Cat. No.98TB100233)*, May 1998, pp. 154–161.

[116] 'The Network Simulator - ns-2'.

[117] N. Papanikos, D.-G. Akestoridis, and E. Papapetrou, 'Adyton: A network simulator for opportunistic networks'. 2015.

[118] nsnam, 'ns-3 a discrete-event network simulator for internet systems', ns-3.

[119] A. Keränen, J. Ott, and T. Kärkkäinen, 'The ONE Simulator for DTN Protocol Evaluation', in *Proceedings of the 2nd International Conference on Simulation Tools and Techniques*, in Simutools '09. ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2009, p. 55:1-55:10.

[120] A. Verma, P. Verma, S. K. Dhurandher, and I. Woungang, *Opportunistic Networks: Fundamentals, Applications and Emerging Trends*. CRC Press, 2021.

[121] M. Liu, Y. Gui, J. Li, and H. Lu, 'Large-Scale Small Satellite Network Simulator: Design and Evaluation', in *2020 3rd International Conference on Hot Information-Centric Networking (HotICN)*, Dec. 2020, pp. 194–199.

[122] F. Neves, A. Cardote, R. Moreira, and S. Sargento, ‘Real-world evaluation of IEEE 802.11p for vehicular networks’, in *Proceedings of the Eighth ACM international workshop on Vehicular inter-networking*, in VANET ’11. New York, NY, USA: Association for Computing Machinery, Sep. 2011, pp. 89–90.

[123] S. Demmel, A. Lambert, D. Gruyer, A. Rakotonirainy, and E. Monacelli, ‘Empirical IEEE 802.11p performance evaluation on test tracks’, in *2012 IEEE Intelligent Vehicles Symposium*, Jun. 2012, pp. 837–842.

[124] Y. Wang, X. Duan, D. Tian, G. Lu, and H. Yu, ‘Throughput and Delay Limits of 802.11p and its Influence on Highway Capacity’, *Procedia - Soc. Behav. Sci.*, vol. 96, pp. 2096–2104, Nov. 2013,

[125] CharIN EV, ‘CCS Specification’.

[126] P. O. of the E. Union, ‘Directive 2014/94/EU of the European Parliament and of the Council of 22 October 2014 on the deployment of alternative fuels infrastructure Text with EEA relevance, CELEX1’.

[127] EV Database (EVDB), ‘Electric Vehicle Database’, EV Database.

[128] SMMT, ‘Motor Industry Facts’, The Society of Motor Manufacturers and Traders, 2022.

[129] EV Database, ‘Tesla Model 3’, Electric Vehicle Database.

[130] Zapmap, ‘How many EV charging points are there in the UK - Zapmap’.

[131] Office for National Statistics, ‘Estimates of the population for the UK, England and Wales, Scotland and Northern Ireland - Office for National Statistics’.

[132] Department for Environment, Food and Rural Affairs, ‘Rural Urban Classification’.

[133] OpenStreetMap contributors, ‘Planet dump retrieved from https://planet.osm.org’. 2021.

[134] N. Basta, A. El-Nahas, H.-P. Grossmann, and S. Abdennadher, ‘Geo-social mobility model for VANET simulation’, *J. Mob. Multimed.*, pp. 107-127-107–127, Jul. 2014.

[135] Department for Transport, ‘Electric vehicle charging device statistics: January 2022’, 2022.

[136] Driver and Vehicle Licensing Agency (DVLA), ‘Vehicle licensing statistics’. Aug. 16, 2023.

[137] Department for Transport, ‘National Travel Survey’, 2020.

[138] G. Falchetta and M. Noussan, ‘Electric vehicle charging network in Europe: An accessibility and deployment trends analysis’, *Transp. Res. Part Transp. Environ.*, vol. 94, p. 102813, May 2021,

[139] F. Bai, N. Sadagopan, and A. Helmy, ‘IMPORTANT: a framework to systematically analyze the Impact of Mobility on Performance of Routing Protocols for Adhoc Networks’, in *IEEE INFOCOM 2003. Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE Cat. No.03CH37428)*, Mar. 2003, pp. 825–835 vol.2.

[140] O. Maghazei, M. A. Lewis, and T. H. Netland, ‘Emerging technologies and the use case: A multi-year study of drone adoption’, *J. Oper. Manag.*, vol. 68, no. 6–7, pp. 560–591, 2022,

[141] M. Torabbeigi, G. J. Lim, and S. J. Kim, ‘Drone Delivery Scheduling Optimization Considering Payload-induced Battery Consumption Rates’, *J. Intell. Robot. Syst.*, vol. 97, no. 3, pp. 471–487, Mar. 2020,

[142] J. Sánchez-García, J. M. García-Campos, M. Arzamendia, D. G. Reina, S. L. Toral, and D. Gregor, 'A survey on unmanned aerial and aquatic vehicle multi-hop networks: Wireless communications, evaluation tools and applications', *Comput. Commun.*, vol. 119, pp. 43–65, Apr. 2018,

[143] J. Jaguemont, L. Boulon, and Y. Dubé, 'A comprehensive review of lithium-ion batteries used in hybrid and electric vehicles at cold temperatures', *Appl. Energy*, vol. 164, pp. 99–114, Feb. 2016,

[144] R. Chertov, S. Fahmy, and N. B. Shroff, 'Fidelity of network simulation and emulation: A case study of TCP-targeted denial of service attacks', *ACM Trans. Model. Comput. Simul.*, vol. 19, no. 1, p. 4:1-4:29, Jan. 2009,

[145] K. Gao *et al.*, 'DONS: Fast and Affordable Discrete Event Network Simulation with Automatic Parallelization', in *Proceedings of the ACM SIGCOMM 2023 Conference*, in ACM SIGCOMM '23. New York, NY, USA: Association for Computing Machinery, Sep. 2023, pp. 167–181.

[146] G. F. Riley, R. M. Fujimoto, and M. H. Ammar, 'A generic framework for parallelization of network simulations', in *MASCOTS '99. Proceedings of the Seventh International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, Oct. 1999, pp. 128–135.

[147] W. Shi, H. Zhou, J. Li, W. Xu, N. Zhang, and X. Shen, 'Drone Assisted Vehicular Networks: Architecture, Challenges and Opportunities', *IEEE Netw.*, vol. 32, no. 3, pp. 130–137, May 2018,

[148] A. Mairaj, A. I. Baba, and A. Y. Javaid, 'Application specific drone simulators: Recent advances and challenges', *Simul. Model. Pract. Theory*, vol. 94, pp. 100–117, Jul. 2019,

[149] M. Radenkovic, A. Walker, and L. Bai, ‘Towards Better Understanding the Challenges of Reliable and Trust-Aware Critical Communications in the Aftermath of Disaster’, in *2018 14th International Wireless Communications Mobile Computing Conference (IWCMC)*, Jun. 2018, pp. 648–653.