

Measuring, Modeling and Enhancing Power-Line Communications

THÈSE N° 7046 (2016)

PRÉSENTÉE LE 17 JUIN 2016

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS

LABORATOIRE POUR LES COMMUNICATIONS INFORMATIQUES ET LEURS APPLICATIONS 3

PROGRAMME DOCTORAL EN INFORMATIQUE ET COMMUNICATIONS

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Christina VLACHOU

acceptée sur proposition du jury:

Prof. J.-Y. Le Boudec, président du jury

Prof. P. Thiran, directeur de thèse

Prof. R. T. B. Ma, rapporteur

Prof. D. Malone, rapporteur

Prof. A. Argyraki, rapporteuse



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2016

As you set out for Ithaka
hope your road is a long one,
full of adventure, full of discovery.
Laistrygonians, Cyclops,
angry Poseidon—don't be afraid of them:
you'll never find things like that on your way
as long as you keep your thoughts raised high,
as long as a rare excitement
stirs your spirit and your body.
...

Ithaka gave you the marvelous journey.
Without her you wouldn't have set out.
She has nothing left to give you now.

And if you find her poor, Ithaka won't have fooled you.
Wise as you will have become, so full of experience,
you'll have understood by then what these Ithakas mean.

— by C. P. Cavafy (translated by Edmund Keeley and Philip Sherrard)

To my family...

Acknowledgements

This thesis was a wonderful journey of about five years, with many ups and downs (as any PhD life is). I am deeply grateful to all the people who made the good moments better and the bad moments not so bad!

I would like to first thank my advisor, Patrick Thiran, for his support, trust, and guidance, and for the knowledge he shared with me during this nice PhD journey. When I met him for the first time, Patrick suggested a project on power-line communications and I was the first to work on this subject in the lab. The beginning of the project was challenging for various reasons. Patrick showed a remarkable understanding of the situation and invested every possible resource in this project. During this challenging research, I was always inspired by the vision and motivation of Patrick. I am deeply grateful to him for hiring me in his lab and for trusting my skills—especially when I did not trust my own skills. I am also grateful for his support of my projects outside EPFL, such as my internships at Qualcomm and Marvell and my job search for my life after my PhD thesis. It is no doubt that Patrick is an exceptional advisor and mentor, a fair person and an amazing teacher.

I would also like to thank the members of my committee, Jean-Yves Le Boudec, Katerina Argyraki, Richard Ma, and David Malone, for their constructive feedback and for the fruitful discussions on my defense. Special thanks to David for his help at the beginning and at the end of my thesis.

This thesis received support and funding from EPFL and the SmartWorld project of the Hasler foundation in Switzerland. I am grateful to both EPFL and Hasler for the grants made available for my salary and trips to conferences, and for systems development. At this point, I would like to thank our lab secretaries Danielle, Angela, Patricia and Holly for making sure that all lab or trip processes run smoothly and for always being kind and helpful. I am very grateful to Holly, not only for “hollyfying” our publications and theses, but also for teaching us how to write correct and meaningful English—skills that we will (hopefully) still use after EPFL. I am also grateful to Yves and Marc-André for their help in our testbed and for the continuous, efficient IT support all these years.

I was very lucky to collaborate with very bright and nice people. I would like to express my gratitude to Albert Banchs, Julien Herzen, Sébastien Henri and Pablo Salvador. Albert has been my mentor and collaborator since 2013, and our collaboration was very fruitful. It was an honor to work with such a leading and bright researcher as Albert. Julien worked with me since the crucial beginning of the project and I would

Acknowledgements

like to thank him for his trust in me, for all the knowledge he shared with me and for his constant constructive feedback. Sébastien arrived later in the lab and completed the “team” of PLC/WiFi networks. I would like to thank him for his help with multiple experiments in this thesis and for his support in times of stress and disappointment. The overall culture in the lab was a determining factor for the success of this thesis; Julien, Vincent, Mohamed, Lucas, Sébastien, Brunella, Matthias, Patrick, Ehsan, Lyudmilla, Young-Jun, Elisa, Emti, Victor and William contributed in the everyday life of the lab. We shared ski and dinner outings, coffee breaks, dry-runs and presentations, paper rejections and acceptances, beers at Sat. This was the “bargest” lab ever! Arriving at the office in the morning was always a joy that I will miss, as the nice, friendly and kind environment of LCA cannot be replaced!

The life in Lausanne would not have been the same without many friends. I would like to express my love to *strobers* (Iris, Sofia, Emre and Lorenzo), who have been my second family in Switzerland. I was lucky to share moments with many nice and fun people: Dorina, Marina, Vicky, Lucia, Giorgia, Matt, Iraklis, Yannis, Manos, Ypatia, Anthi, Panagiotis, Nadia, Vassilis T., and last but not least Vassilis K. (with whom I discovered all hidden bars of Lausanne). Thank you all for your support, trips, endless discussions, Greek traditions and good times. Thank you, Evgenia, Mairy, Evita, Maria, Ioanna, for your distant support from Greece, for all the nice moments before and during my PhD thesis.

I would like to especially thank a few people who only came into my life at the last year of this journey, but their support was critical at the last steps, such as the defense and the submission of this thesis. In particular, I would like to thank Can for the wonderful moments in multiple places and trips on this planet, for his support during the thesis and job-search processes, for making me happy. I would also like to thank my flatmate Dimitra—who arrived home at the very last steps of this journey—for her support and her “bonne humeur”!

Finally, I would like to express my love and gratitude to my family. I have adopted various qualities from them and I am really proud to be part of such a family. My father was the one who initiated me to computers and motivated me to study computer engineering! I am always inspired by his integrity, fairness, altruism and positive attitude. I would like to thank my mother for standing by me, taking care of me even in distance and for teaching me to be strong (even in cases of failure). I would like to thank my sister Penny for all the wonderful moments and support. I am very lucky to be part of such a nice and supportive family; Everything I Am, I Owe To You!

Lausanne, 17 June 2016

Christina Vlachou

Abstract

Power-line communication (PLC) is a technology that has become increasingly popular in recent years, as it provides easy and high-throughput connectivity in residential and enterprise networks. The IEEE has standardized PLC with the 1901 standard, and the HomePlug Alliance certifies products and provides specifications for different data-rates and applications. Furthermore, tens of vendors offer hybrid solutions with both WiFi and PLC, thus yielding an extended coverage and relieving congestion in the network.

Although it is commercially successful and widely adopted, PLC has received far too little attention from the research community. Two of the main reasons for this lack of research are the proprietary nature of the technology and the protocols' complexity. We overcome these barriers and build the foundations for understanding, evaluating, measuring, exploiting and boosting PLC performance. In this dissertation, we explore PLC performance and dynamics in various time-scales and treat both best-effort and delay-sensitive applications towards enhancing quality of service.

When deploying hybrid networks, there are two open questions that arise: Does PLC perform better than WiFi and to which extent can PLC augment network reliability? How can we accurately estimate PLC capacity for deciding to which medium data should be forwarded? To answer these questions, we conduct an experimental study with PLC and WiFi stations and delve into the spatio-temporal variations of PLC capacity. Our results uncover crucial differences between the two mediums, thus proving that PLC largely extends coverage and augments network reliability. We discover that PLC links are strongly asymmetric and that temporal variation occurs on three time-scales. There is a high correlation between link quality and its variability, which has a direct impact on probing overhead and on accurate link-metric estimations. Not only do we make it possible for the first time to understand the key features of PLC, but we also propose systematic guidelines for PLC link-metric estimation.

The subsequent open questions are related to the efficiency of PLC, when stations contend for the medium. We investigate the PLC MAC layer and uncover that PLC employs a CSMA/CA mechanism that avoids wasting time for collisions, due to PHY-layer features, such as the significantly large frame duration with respect to the time-slot duration. This CSMA/CA bears a resemblance to that of WiFi, in the sense that it includes a binary exponential backoff. WiFi stations double the contention window only after experiencing a collision. In contrast, PLC enables the stations to also double their contention window before experiencing a collision. To this end, PLC introduces

an additional variable, called the deferral counter, that regulates the frequency of this proactive reaction based on congestion in the network. We introduce a model for evaluating performance. The model relies on the so-called decoupling assumption that asserts that the backoff processes of the stations are independent and has been widely used for modeling WiFi networks. Despite the deferral counter intricacies, our model boils down to a single fixed-point equation of the collision probability. We prove the uniqueness of the solution for a wide range of configurations and also exploit the model to devise configurations that significantly boost performance for best-effort applications. We corroborate our model and performance gains via extensive simulation and measurements on WiFi and PLC hardware.

After showing that PLC CSMA/CA is more efficient than that of WiFi, we ask the question: At which price does this efficiency come? We now turn our focus to short-scale dynamics. We unveil analytically, experimentally and in simulation that, contrary to WiFi, PLC is short-term unfair. This yields high delay-variance, that is, jitter, and it affects delay-sensitive applications. The deferral counter introduces unfairness and also determines a tradeoff between throughput and fairness, which is extensively studied. We also find that the PLC unfairness leads to strong dependence between stations, which can penalize the accuracy of models relying on the decoupling assumption.

To further improve the modeling accuracy of PLC CSMA/CA, we propose another model that does not resort to the decoupling assumption. In this coupled model, we allow the transition probabilities of a tagged station to depend on the backoff processes of the other stations, thus taking into account the short-scale dynamics of PLC. The resulting model, which is a dynamical system, is more complex compared to the first one, but it performs better for small number of stations, a frequent scenario in practice. Here too, we prove that the model admits a unique steady-state solution for a wide range of configurations. This is the first model of the PLC CSMA/CA that reaches this level of accuracy. By using this model and our research on fairness, we propose an algorithm that yields configurations with low jitter, given throughput constraints and a set of possible configurations.

Keywords: Power-line communications, HomePlug, IEEE 1901, spatio-temporal variations, capacity estimation, CSMA/CA, deferral counter, performance evaluation, decoupling assumption, fairness.

Résumé

La communication par courants porteurs en ligne (CPL) est une technologie qui devient de plus en plus populaire, parce qu'elle offre une connexion haut débit facile à installer pour les réseaux locaux. IEEE a standardisé les CPL avec le standard 1901, et l'alliance HomePlug certifie des produits et offre des spécifications pour des débits et des applications diverses. De plus, des dizaines de vendeurs offrent des solutions hybrides avec WiFi et CPL, rendant possible une couverture étendue et allégeant la congestion dans le réseau.

Malgré sa large adoption et son succès commercial, les CPL ont trop peu reçu l'attention du milieu de la recherche. Les deux raisons principales de ce désintérêt sont le caractère propriétaire de cette technologie et la complexité des protocoles. Nous surmontons ces barrières et nous jetons les bases pour comprendre, mesurer, exploiter et améliorer la performance des CPL. Dans cette thèse, nous examinons la performance et la dynamique des CPL à plusieurs échelles de temps et nous nous intéressons à des applications qui sont sensibles au débit ou au délai, améliorant donc la qualité de service.

Deux questions liées au déploiement des réseaux hybrides restent ouvertes. Est-ce que les CPL fonctionnent mieux que le WiFi et jusqu'à quel point les CPL peuvent améliorer la fiabilité du réseau? Comment peut-on estimer précisément la capacité des CPL pour décider vers quelle technologie envoyer les données? Pour répondre à ces questions, nous conduisons une étude expérimentale avec des stations CPL et WiFi, et nous approfondissons la variabilité spatiale et temporelle de la capacité des CPL. Nos résultats montrent des différences cruciales entre les deux technologies et prouvent que les CPL étendent largement la couverture et améliorent la fiabilité du réseau. Nous découvrons que les liens CPL sont très asymétriques et que des variations temporelles existent à trois échelles de temps. Il y a une forte corrélation entre la qualité du lien et sa variabilité, ce qui a une influence directe sur la précision et le coût de l'estimation des métriques de qualité de lien. Pour la première fois, nous permettons de comprendre les principales caractéristiques des CPL, et nous proposons également des lignes directrices systématiques pour l'estimation des métriques CPL.

Les questions suivantes sont liées à l'efficacité des CPL lorsque plusieurs stations essaient d'accéder au support. Nous examinons la couche d'accès (MAC) des CPL et nous trouvons que les CPL utilisent un mécanisme CSMA/CA qui évite une perte de temps due aux collisions, à cause des particularités de la couche physique (PHY), comme la durée longue du paquet par rapport à l'intervalle de temps. Ce CSMA/CA est similaire à celui du WiFi en ce qu'il emploie un backoff exponentiel binaire. En revanche, les stations WiFi

doublent leur fenêtre de contention (*contention window*) seulement après une collision, alors que les stations CPL peuvent également doubler leur fenêtre de contention avant une collision. Pour cela, le standard CPL introduit une variable nouvelle, appelée le compteur d'échec d'émission (*deferral counter*), qui régit la fréquence de cette réaction proactive fondée sur la congestion dans le réseau. Nous introduisons un modèle pour évaluer la performance du CSMA/CA des CPL. Le modèle repose sur une hypothèse de découplage qui signifie que les processus de backoff des stations sont indépendants, hypothèse qui est largement utilisée pour l'analyse des réseaux WiFi. Malgré la complexité du compteur d'échec d'émission, notre modèle se réduit à une équation de point-fixe fonction de la probabilité de collision. Nous prouvons qu'il y a une solution unique pour une vaste gamme de configurations, et nous exploitons également le modèle afin d'élaborer des configurations qui améliorent largement la performance. Nous corroborons notre modèle et nos bénéfices de performance par des simulations approfondies et des expériences sur du matériel WiFi et CPL.

Après avoir découvert que le CSMA/CA des CPL est plus efficace que celui du WiFi, nous nous posons la question suivante : quel est le compromis à trouver pour cette efficacité ? Nous nous concentrons donc ensuite sur les dynamiques à court terme. Nous montrons par analyse, expériences et simulations que, contrairement au WiFi, les CPL sont inévitables à court terme. Cette propriété induit de fortes variations du délai (*jitter*) qui affecte certaines applications sensibles au délai. Le compteur d'échec d'émission introduit une certaine iniquité, mais permet également d'augmenter les performances ; nous étudions largement ce compromis entre performance et équité. Nous trouvons que l'iniquité cause un couplage fort entre les stations, qui peut pénaliser la précision des modèles fondés sur l'hypothèse de découplage.

Afin d'améliorer la précision de l'analyse du CSMA/CA des CPL, nous proposons un autre modèle qui n'est pas fondé sur l'hypothèse de découplage. Avec ce modèle, nous autorisons les probabilités de transition de chaque station à dépendre du processus de backoff d'autres stations, prenant aussi en compte les dynamiques à court terme des CPL. Le modèle obtenu, qui est un système dynamique, est plus complexe que notre premier modèle, mais il accomplit de meilleures performances avec des petites valeurs du nombre de stations, scénario très fréquent. Nous prouvons que ce modèle a une solution unique pour une vaste gamme des configurations. Il s'agit du premier modèle du CSMA/CA des CPL qui obtient ce niveau de précision. En utilisant ce modèle et nos recherches sur l'équité de la couche MAC, nous proposons un algorithme qui retourne des configurations avec une faible variance du délai, étant données une contrainte sur le débit et une gamme de configurations possibles.

Mots clefs : Communication par courants porteurs en ligne, HomePlug, IEEE 1901, variation spatiale et temporelle, estimation de la capacité, CSMA/CA, le compteur d'échec d'émission, évaluation des performances, hypothèse de découplage, équité.

List of Abbreviations

AP	Access point
ARP	Address resolution protocol
ARQ	Automatic repeat request protocol
BLE	Bit loading estimate
BPSK	Binary Phase Shift Keying
CCo	Central coordinator
CSMA/CA	Carrier sense multiple access with collision avoidance
ETX	Expected transmission count
ETT	Expected transmission time
FEC	Forward error correction
FFT	Fast Fourier transform
HPAV	HomePlug AV
MAC	Medium access control
MCS	Modulation and coding scheme index
MIMO	Multiple input multiple output
MM	Management message
MPDU	MAC protocol data unit
NEK	Network encryption key
NMK	Network membership key
OFDM	Orthogonal frequency division multiplexing
PB	Physical block
PIB	Platform independent binary
PLC	Power-line communications
QAM	Quadrature amplitude modulation
QPSK	Quadrature phase shift keying
RTS/CTS	Request to send/Clear to send
SACK	Selective acknowledgement
SoF	Start of frame delimiter
SWM	Sliding window method
TCC	Turbo convolutional coder
TCP	Transport control protocol
TDMA	Time division multiple access
TEI	Terminal equipment identifier
UDP	User datagram protocol

Contents

Acknowledgements	i
Abstract (English/Français)	iii
List of Abbreviations	vii
1 Introduction	1
1.1 Today's Networks: Connectivity Problems and Demands	1
1.2 Augmenting Reliability with Hybrid Networks	3
1.2.1 Hybrid Standardization and Solutions	4
1.2.2 Candidate Technologies for Hybrid Networks	4
1.2.3 Power-Line Communication: Easy, High-Throughput Connectivity	5
1.3 PLC Research: Challenges and Open Problems	6
1.4 Dissertation Outline and Contributions	10
1.4.1 WiFi vs. PLC: Which One Performs Better?	13
2 A Crash Course on Power-Line Communication	15
2.1 Introduction	15
2.2 Channel Characteristics	16
2.2.1 Attenuation and Frequency Selectivity	16
2.2.2 Types of Noise	17
2.3 Standardization and Specifications	18
2.4 PHY Layer	19
2.4.1 Modulation and Coding	19
2.4.2 Channel Estimation and Adaptation	20
2.5 MAC Layer	21
2.5.1 Frame Structure and Aggregation	22
2.5.2 Access Control: The IEEE 1901 CSMA/CA Protocol	23
2.6 Network Structure and Management	27
2.7 WiFi vs. PLC	28
2.7.1 Spatial Variation	29
2.7.2 Temporal Variation	31
2.8 Summary	32

3	Spatio-temporal Variations and Estimation of Capacity	33
3.1	Introduction	33
3.2	Spatial Variation of PLC	34
3.2.1	Topology and Connectivity	34
3.2.2	Asymmetry in PLC Links and Consequences	35
3.3	Temporal Variation of PLC	37
3.3.1	Invariance Scale	38
3.3.2	Cycle Scale	39
3.3.3	Random Scale	42
3.4	Capacity-Estimation Process	44
3.4.1	The BLE as a Capacity Estimator	44
3.4.2	Size of Probe Packets	45
3.4.3	Priority Class of Probe and Control Packets	46
3.4.4	Validation of our Capacity Estimation in Hybrid Single-Hop	47
3.5	Retransmitting in PLC Channels	49
3.5.1	Retransmission Due to Errors	49
3.5.2	Retransmission Due to Contention	51
3.6	Link-Metric Guidelines	53
3.7	Related Work	53
3.8	Summary	54
4	Decoupled Analysis and Enhancement of Multi-user Performance	57
4.1	Introduction	57
4.1.1	Background and Motivation: IEEE 802.11 vs. 1901 CSMA/CA	57
4.2	Analysis	60
4.2.1	Network Assumptions	60
4.2.2	Decoupling Assumption	61
4.2.3	The Decoupling-Assumption Model	61
4.3	Enhancements of IEEE 1901 for High Throughput	64
4.3.1	Deriving the Proposed Configuration	64
4.3.2	1901-Proposed Advantage over 802.11	68
4.4	Performance Evaluation	69
4.4.1	Experimental Validation of our Simulator and Model	69
4.4.2	Simulation Assumptions and Parameters	70
4.4.3	Validation of the Decoupling-Assumption Model	71
4.4.4	Proposed Enhancement for the IEEE 1901 MAC	71
4.4.5	Experimental Validation of Proposed Enhancement	75
4.5	Related Work	77
4.6	Summary	79

5	Short-Term Fairness and Dynamics in Multi-user Settings	81
5.1	Introduction	81
5.2	Metrics of Fairness	83
5.3	Analysis	84
5.3.1	Assumptions	85
5.3.2	The Distribution of Inter-transmissions	85
5.4	Fairness Evaluation via Simulation and Testbed	88
5.5	Short-Term Fairness and Impact on Delay	91
5.6	The Tradeoff between Throughput and Short-Term Fairness	92
5.7	Fairness and Coupling between the Stations	95
5.8	Related Work	97
5.9	Summary	97
6	Coupled Analysis of Multi-user Performance	99
6.1	Introduction	99
6.2	Analysis	100
6.2.1	The Drift Model	100
6.2.2	Uniqueness of the Equilibrium Point	105
6.2.3	Convergence to the Unique Equilibrium Point	106
6.3	Performance Evaluation	108
6.3.1	Drift vs. Decoupling-Assumption Model	109
6.3.2	Uniqueness of the Solution and Selected Counterexample	111
6.3.3	Accuracy of the Deterministic System in Transient Regime	113
6.4	Adjusting the Fairness-Throughput Tradeoff with the Drift Model	114
6.5	Related Work and Discussion	117
6.6	Summary	118
7	Conclusion	121
A	Proofs	123
A.1	Decoupling-Assumption Model	124
A.2	Drift Model	130
A.3	Distribution of Inter-transmissions	138
B	Experimental Framework	143
B.1	Testbed and Setup	143
B.2	Measurement and Traffic Tools	144
B.3	MAC-Layer Experiments	145
	Bibliography	151
	Curriculum Vitae	153

1 Introduction

1.1 Today's Networks: Connectivity Problems and Demands

As we move to the era of connected homes and a myriad of mobile applications, communications need to be scaled accordingly and to provide reliable, full coverage. Although wireless communications are pervasive and their efficiency has been largely improved by recent amendments, there are still coverage problems and the network often reaches high levels of congestion, thus failing to guarantee high quality-of-service. Nowadays, we observe an explosion of connected possibilities in all kinds of environments, from residential and enterprise to rural and urban, which is a challenge for state-of-the-art networking solutions. In addition, users have extremely stringent requirements from communication technologies, such as mobility, ease-of-use, performance, and reliability.

Wireless technology is dominant in residential and enterprise networks; it offers mobility and attractive data-rates. Nevertheless, it often leaves “*blind spots*” in coverage due to building construction. Wireless signals are severely attenuated by walls built of brick or stucco that act as an obstacle to Internet access for certain users. A recent study in 1000 American households suggests that 40% of residences with wireless routers have experienced problems with their network [1]. Among those households experiencing connectivity problems, 63% continue to have these problems despite efforts to resolve them. A solution often employed incorporates wireless repeaters and multi-hop communication. Yet, with this topology, wireless signals are still obstructed by construction obstacles, often failing to reliably extend coverage.

Due to the shared nature of wireless medium, another connectivity issue arises due to *neighboring networks that cause interference* to each other. The situation is usually alleviated by operating at a different frequency channel, non-interfering with neighboring networks. However, channel allocation is becoming increasingly challenging: due first, to the density of wireless networks in urban environments and second, to recent WiFi standardizations, such as IEEE 802.11ac, that maximize utilized bandwidth yielding only one or two available orthogonal channels.

In addition to wireless signal-attenuation and interference, there are connectivity intricacies with wired technologies that are bounded to the connected- or smart-devices. For instance, smart devices, such as TVs, often enable only Ethernet connectivity, which requires *costly infrastructure*. Even when smart devices offer WiFi capabilities, the users often prefer wired connections for non-mobile applications in order to avoid wasting WiFi bandwidth. Obviously, compared to WiFi, Ethernet cannot have ubiquitous coverage and the ease-of-use is limited due to long-wiring and high cost.

Wireless networking problems are exacerbated by the congestion of connected devices. In recent years, a vast amount of mobile applications has been added in our every day life: smart-phones and tablets, wearables, home-automation devices for comfort, energy savings, security. In fact, Cisco and Goldman Sachs estimate that by 2020, we will have more than 50 billion connected devices and that more than 250 new “things” will be connected each second [2, 3]. The explosion in the number of connected devices per building and the “Internet-of-things” (IoT) networking require *communication technologies and protocols that can scale and that react efficiently to congestion*.

Beyond IoT and smart-devices, today the decrease in bandwidth and processing costs enables many bandwidth-hungry applications, such as augmented-reality, ultra high-definition streaming, gaming, video and web conferencing, cloud storage, virtual desktops. Clearly, WiFi fails to meet the required quality of service in case of multiple applications or contending flows, despite the recent amendments for high efficiency. These amendments have pushed most WiFi parameters to the capacity limit (e.g., channel bandwidth, modulation, number of antennas and spatial streams). Hence, in the future, WiFi capacity improvements will be limited. Moreover, as mentioned earlier, wired Ethernet technologies—often the most popular substitute of WiFi—require costly infrastructure, confining their ubiquitous coverage. For all these reasons, *additional communication mediums are required for bandwidth-hungry applications*.

The aforementioned connectivity problems and demands call for *hybrid networks* combining multiple efficient and user-friendly communication technologies. The incorporation of additional communication technologies in today’s networks can bring multiple benefits. In addition to bandwidth aggregation, diverse mediums can augment network reliability, extend coverage and provide different levels of quality-of-service in the network. Harnessing multiple technologies enables exploitation of the medium diversity and accommodation of more users and applications. This will enable the user to eliminate connectivity problems and meet high-rate demands, and to have additional gains such as resilience, reliable backbones, seamless connectivity.

In this dissertation, we explore a promising communication technology for hybrid residential and enterprise networks: power-line communication (PLC). PLC is beneficial in terms of both easy installation and high data-rates: It enables data transmission via the electrical wires—an ubiquitous infrastructure—, it is trivial to install, and it provides rates up to 1.5 Gbps. Most common applications of PLC include WiFi coverage extension and replacement of costly Ethernet connections, depicted in Figure 1.1. Although PLC is




Figure 1.1 – Examples of PLC addressing connectivity problems and demands.

increasingly popular, it has received far too little attention from the research community, especially from an end-user performance perspective. The proprietary nature of the technology and the extreme protocol complexity are the main reasons for this lack of research. Overcoming these challenges, we present a thorough evaluation of PLC under both single- and multi-user settings: We resolve fundamental open problems on performance, we introduce performance enhancements, and we facilitate PLC future integration into hybrid networks.

In the remainder of this chapter, we first explore the evolution of hybrid networks and candidate technologies for augmenting network reliability. We then turn our attention to power-line communications and we highlight the reasons it is promising, the research challenges and the open problems around this technology. Finally, we present an outline of this dissertation along with our contributions.

1.2 Augmenting Reliability with Hybrid Networks

We now discuss the state-of-the-art efforts for combining multiple communication technologies towards hybrid networks. WiFi is always included in hybrid solutions, as it offers mobility and serves the vast majority of mobile applications. We explore candidate technologies that have the potential to augment WiFi reliability, and we discuss their perks and deficiencies.

1.2.1 Hybrid Standardization and Solutions

As the demand for combining diverse, non-interfering technologies increases, new specifications for hybrid networks have been developed, such as the IEEE 1905 standard [4] which specifies abstraction layers for topology, link metrics, and forwarding rules. Hybrid networks operate at layer 2.5, i.e., between IP and MAC layers, according to standardization and commercial solutions. In this way, they provide seamless connectivity and interoperability between different mediums.

IEEE 1905 provides a common interface for widely deployed home networking technologies, such as wireless, power-line, coaxial and Ethernet communications. The nVoy program [5] is a certification and marketing initiative implementing the IEEE 1905 hybrid networking standard. Tens of vendors have released hybrid products including PLC and WiFi, or coaxial-communications and WiFi and collaborate with the nVoy alliance towards certifications of their solutions. In addition to reliability, full coverage and bandwidth aggregation, nVoy targets energy management, simple and push-button security, and advanced network-diagnostics.

1.2.2 Candidate Technologies for Hybrid Networks

As mentioned in the previous section, several candidates, among which power-line and coaxial communications, for ameliorating WiFi performance are on the market. We now review these technologies in the quest for reliable, cost-effective and easy-to-use solutions. Table 1.1 summarizes the value proposition¹ and drawbacks of four popular communication technologies. Let us discuss each technology separately.

WiFi is undoubtedly the most popular solution due to the explosion of mobile applications. As WiFi popularity is expected to grow, the other candidate technologies usually act as a backbone for WiFi. The main reasons for the existence of these backbones are the connectivity demands and problems discussed in Section 1.1. In recent years, multi-channel WiFi is widely employed, with multiple WiFi devices simultaneously utilizing channels at 2.4, 5 and 60 GHz bands. The advantage of such solutions is clearly mainly bandwidth aggregation; coverage extension and resilience cannot be achieved, as all wireless-bands capacities are affected by the same or correlated factors (e.g., signal penetration, fading).

PLC is becoming very popular in home networks. The main perk of this technology is the no-new-wires connectivity and the high density of electrical plugs in any residential or enterprise environment. Furthermore, PLC is a natural and trivial solution for smart electrical appliances already connected to the grid. The potential disadvantages of PLC are interference—because PLC is a shared medium like WiFi—and high performance variability depending on the age of the electrical grid or the electrical appliances operating.

¹“A value proposition is a business or marketing statement that summarizes why a consumer should buy a product or use a service” [6].

1.2. Augmenting Reliability with Hybrid Networks





Technology	Value proposition	Drawbacks
WiFi 	Mobility and ease of use	<ul style="list-style-type: none"> • Prone to interference • Too many devices on the network affect usable bandwidth • Walls cause severe attenuation
PLC 	High availability of outlets and ease of use	<ul style="list-style-type: none"> • Prone to interference • Too many devices on the network affect usable bandwidth • Performance varies depending on wiring quality and electrical-grid structure
MoCA 	High performance and reliability	Dependent on number of coaxial outlets
Ethernet 	High performance and reliability	Requires new wiring and costly infrastructure

Table 1.1 – IEEE 1905 candidate technologies for hybrid networks: value proposition and drawbacks. Source of data: MoCA alliance [7].

Coaxial communication is specified and designed by the Multimedia over Coax Alliance (MoCA) [7], and it is usually known simply as MoCA. Communication over coaxial cables achieves 1.4 Gbps data-rates, employing bands in the frequency range of 500–1650 MHz and with 100 MHz of bandwidth. The main advantage of this technology is high reliability with packet error probability of about 10^{-6} . The interference is also limited within local network, as shielded coaxial wires prevent interference from neighboring households or other technologies. MoCA applications mainly focus on multi-room digital video recorders, over-the-top streaming content, gaming, ultra HD streaming. Although it provides high reliability and rates, MoCA requires an already established coaxial system and possibly more than two coaxial plugs that are limited in today’s households.

Finally, Ethernet communication provides the highest reliability, but is the most costly. It is usually deployed in enterprise buildings with a sophisticated network structure. In residential environments, users often employ Ethernet for non-mobile applications, where the main disadvantage is the requirement for new wiring.

Therefore, PLC is certainly a promising technology in terms of cost, flexibility and pervasiveness of connectivity possibilities. In the following section, we give more details and explain advantages of this technology.

1.2.3 Power-Line Communication: Easy, High-Throughput Connectivity

Due to the growing demand of reliability in home networks, wireless and power-line communications are combined by several vendors to deliver high rates and broad coverage

without blind spots. PLC is at the forefront of home networking, as it provides easy, plug-n-play, and high data-rate connectivity. Its main advantage is coverage wider than WiFi and high data-rates up to 1.5 Gbps without requiring the wiring of a new network.

PLC can certainly handle bandwidth-hungry applications, but it can also be employed for home automation and low-rate networking. Electrical smart-appliances are inherently connected to the power-grid, hence they can use the power-line for both energy-supply and communications. This property of PLC is very efficient for future connected homes, as they will use multiple smart appliances for comfort, security and energy savings.

Because of all advantages discussed above, PLC is widely adopted in residential and enterprise networks. HomePlug [8], the leading alliance for PLC standardization and certifications, estimates that over 180 million PLC devices have already been shipped, and that the expected annual growth will be more than 31% the next years [9]. HomePlug also proposes different solutions for home automation and high data-rate local networks. The most popular specification for high data-rate PLC, employed by 95% of PLC devices [9], is HomePlug AV. This specification was adopted by the IEEE 1901 standard [10].

In this dissertation, we focus on indoor and broadband PLC, in the frequency range of 1–80 MHz, and more specifically on the HomePlug AV and IEEE 1901 specifications. However, note that there also exist narrowband solutions for outdoor reliable communications with applications on power-distribution automation, demand-response control, meter-to-grid connectivity. These solutions are implemented on rural, urban, and suburban areas. The IEEE 1901.2 standard [11] provides guidelines for narrowband PLC using frequencies below 500 kHz and data rates of up to 500 kbps. Narrowband PLC technologies are out of the scope of this work.

1.3 PLC Research: Challenges and Open Problems

Power-line communication is developing rapidly. Although it is commercially popular and widely adopted, this technology has received far too little attention from the research community. The vast majority of publications in PLC focus on channel modeling and little work has been done on end-to-end performance from a user perspective. In this section, we first uncover the challenges of studying PLC, which are most probably the reasons for this lack of research. Then, we describe the main open problems in PLC that this thesis solves. The main *challenges* to studying PLC are the following.

Intricate channel characteristics. Power-lines were not designed for communication purposes. Hence, there are multiple grid-components that create high channel attenuation. Moreover, the main sources of noise in PLC, which are electrical appliances, have unique noise characteristics both in frequency and time. These intricacies challenge end-to-end performance modeling and accurate capacity prediction at different environments.

Complex PHY/MAC design and no justification for intricacies. The specifica-

1.3. PLC Research: Challenges and Open Problems

tions that have dealt with the harsh PLC channel conditions yield a very complex PHY and MAC design, especially compared to the one of WiFi. This complexity challenges the simulation and modeling of PLC performance. Furthermore, the complexity is often not justified by any standardization or publication thus, requiring intensive reverse engineering in order to understand performance.

Proprietary protocols and encrypted firmware. PLC standardization and specifications leave many implementation and design details unspecified. The vendor is responsible for determining the protocol and parameter values of unspecified processes in the standard. Due to the proprietary nature of the technology, the firmware of state-of-the-art PLC solutions is encrypted and the access to various vendor-specific mechanisms, protocols, and parameters is not possible.

In this dissertation, we overcome the barriers described above. By doing so, we address fundamental research questions about PLC. There are numerous *open problems* on PLC performance. These are both relative to end-user performance and to a hybrid-networks design that fully exploits PLC. We now discuss four of these problems and their implications.

Justification and Corroboration of PLC Complexities

One of the open questions about PLC is why actual implementations are so intricate. As we will see later, both PHY and MAC layers are more complex than those of WiFi. Although usually this complexity offers more potential compared to WiFi, until now, there has not been any justification for the design choices and the values of protocol parameters. Explaining the key PLC parameters could ease future research and deployment.

Corroborating protocols and implementations is important not only for PLC, but also for hybrid networks. The PLC intricacies and the lack of experimental frameworks hinder the design, research and implementation of hybrid networks. Hybrid networks require a solid understanding of the underlying layers of each network technology in order to fully exploit each medium.

Capacity Estimation and Variability in PLC and Hybrid Networks

As we will observe later, although PLC boosts network performance, there might be still a few links that perform poorly with both WiFi and PLC. As a result, mesh configurations (hence routing and load balancing algorithms) are needed for seamless connectivity in home or office environments. A challenge for these algorithms is that they have to deal with two different interference graphs with diverse spatio-temporal variation, and that, to fully exploit all mediums, they require accurate metrics for capacity and packet loss-rates. The IEEE 1905 standard specifies link metrics for hybrid networks. However, this standard is technology agnostic and it does not provide any forwarding or metric-estimation methods for the candidate technologies.

A significant challenge, highlighted by recent studies in 802.11n networks [12], is the accuracy of well-established link-quality metrics, such as the expected transmission count (ETX) or time (ETT) [13], in modern networks, i.e., the IEEE 802.11n/ac for WiFi and high-rate PLC. Sheshadri and Koutsonikolas [12] show that due to the MAC/PHY enhancements of 802.11n, these metrics perform poorly and that they should be revised, given that they have been evaluated only under 802.11a/b/g and that they employ mainly broadcast probes. To this end, unicast probes can be exchanged among the stations². However, in a network of N stations, unicast probing introduces an $\mathcal{O}(N^2)$ overhead. The probing overhead can be significantly reduced without significant accuracy loss by employing temporal variation studies of each medium and by probing less often links with low temporal variability. Overall, two research questions, with respect to link-metrics in hybrid networks are (i) the evaluation of well-established link metrics, and (ii) the design of new link metrics and their optimization with respect to both overhead and accuracy.

From both hybrid-networks and a user perspective, there are two research directions to follow for PLC performance variability:

Spatial variation: (i) Which locations ensure good PLC links? (ii) What does the PLC topology look like? (iii) How should PLC networks be designed in large-scale deployments?

Temporal variation: (i) How stable are PLC links with respect to time? (ii) How often does the capacity change and at which variance? (iii) How can probing and link-metric estimation be designed to achieve low overhead and high accuracy?

Because PLC is expected to be a lucrative backbone for WiFi, more questions are raised: Where and when does PLC perform better than WiFi? What are the differences between the two technologies and which medium(s) should an application use?

To conclude, the main questions about PLC performance described above are the (i) link-quality and capacity estimation techniques; (ii) temporal and spatial variation studies; (iii) performance comparison with WiFi.

PLC Multi-user Performance Evaluation

PLC is a shared medium like WiFi, hence it employs a MAC layer tailored for resource allocation. Even though the vast majority of PLC devices follow the HomePlug specification, the MAC layer has received little attention so far from the research community. In particular, no work has investigated how far from optimality this MAC protocol is.

As PLC technology is becoming an important component in home networks, residential buildings are expected to host networks with a high number of PLC stations. These PLC stations interfere with each other, because—in contrast to wireless technologies that rely on different communication channels—PLC utilizes the entire available bandwidth (1.8–80 MHz) for communication. Therefore, there is a need for enhancements at the

²Broadcast packets cannot be used to accurately estimate capacity. See for example [13, 14].

1.3. PLC Research: Challenges and Open Problems

MAC layer, as an efficient MAC is essential to maintain good performance when many stations contend for the medium.

Due to the shared nature of power lines, HomePlug devices employ a multiple-access scheme based on CSMA/CA, which is specified by the IEEE 1901 standard. The 1901 CSMA/CA protocol bears some resemblance to the CSMA/CA mechanism employed by IEEE 802.11, which has been extensively studied in the literature. Both technologies use a *binary exponential backoff* procedure according to which the station selects an interval, called the backoff counter, uniformly at random between 0 and the contention window; it then defers for backoff-counter idle slots before transmitting in order to avoid collisions. In case of a collision, the station doubles the contention window and attempts retransmission with a new backoff counter. Nevertheless, 1901 differs from 802.11 in that its CSMA/CA mechanism is more complex, making its theoretical analysis challenging. In particular, in addition to using a backoff counter, PLC also uses a so-called *deferral counter*. The deferral counter significantly increases the state-space required to describe the backoff procedure, which contrasts with the comparatively small state-space required to analyze 802.11 (see, e.g., the Markov chain used in [15]). As a result, the analysis of 1901 has received little attention despite the commercial success and massive adoption of PLC technologies.

From a general perspective, it turns out that 1901 implements an approach to contention resolution that differs drastically from the 802.11 CSMA/CA. In particular, whereas 802.11 can react to only contention (by doubling its contention window and reducing its transmission probability) after detecting a collision, 1901 can proactively react when it senses the medium busy for a certain number of time slots (given by the deferral counter). Such a protocol design has two distinct advantages over 802.11:

1. The contention window can be increased as many times as required to reach appropriate backoff durations without suffering any collision. In contrast, with 802.11 the contention window can only be doubled after a collision, thus in many cases several collisions need to occur before the contention window reaches the appropriate value. As a result, 1901 can reduce the channel time wasted in collisions, potentially leading to better performance.
2. By appropriately selecting the number of busy slots that trigger an increase of the contention window, we can adjust with fine granularity the level of contention that triggers a reaction. In contrast, this is not possible in 802.11, where contention is detected by the binary signal given by channel occupation: either the channel is busy upon a transmission attempt, which yields a collision, or it is not, and any finer refinements are not possible.

The above reasoning suggests that 1901 can substantially outperform 802.11 if properly configured. However, as we will observe, the default configuration of 1901 does not achieve the level of efficiency that would be expected given these premises. One important cause of the (relatively) poor performance of the protocol is the lack of an accurate and simple analysis that provides an insightful understanding of its dynamics

and that can be used to configure the protocol appropriately. The performance evaluation of 1901 CSMA/CA mechanism is interesting not only from a PLC aspect, but also from a general access-control protocols point of view.

This lack of simple analysis for the 1901 CSMA/CA has further implications. Until now, it has not been proven whether the model of 1901 yields a unique solution at steady-state. The proof of uniqueness to CSMA/CA models is crucial not only from a modeling perspective, but also from a performance one. Configurations with more than one stationary regime have been shown to cause extreme situations of unfairness and multi-stability for 802.11 [16]. Unfairness can starve certain stations. Hence, an open problem in PLC multi-user analysis is the existence of unique stationary regime.

Short-Term Performance Dynamics in PLC Multi-user Settings

Going beyond PLC average performance in steady-state, it is interesting to investigate the short-term dynamics and the ability of the MAC protocol to allocate efficiently and in a fair way the available resources to contending stations. In particular, if fairness is not guaranteed among stations, then substantial delay-variance is introduced. Therefore, studies of short-term dynamics are crucial for delay-sensitive applications that have low-jitter requirements. Compared to the extensively studied WiFi, there are no studies on the short-term dynamics of the PLC CSMA/CA.

Short-term dynamics are much more important with PLC than WiFi due to the introduction of deferral counter. Note that this counter creates strong coupling between the stations, because it regulates the frequency of the reaction to sensed transmissions in the channel. Deferral-counter dynamics infers a different short-term behavior compared to WiFi. Moreover, it challenges the accuracy of modeling approaches that assume decoupling (independence) between the stations and that have been well established and justified for WiFi. By delving into the short-term dynamics, we can gain insights into both performance and rigorous modeling of multi-user scenarios.

1.4 Dissertation Outline and Contributions

As we discussed above, there are fundamental open questions about PLC performance. This dissertation addresses these aforementioned open problems. We now present the outline of this dissertation and highlight the main contributions of each chapter.

Chapter 2 We present a tutorial on the main characteristics and features of power-line communications. We discuss the following: *(i)* main characteristics of power-line channel, such as noise and attenuation, *(ii)* most popular specifications and their differences, *(iii)* PHY and MAC layers in practice, and *(iv)* network structure and management techniques. By justifying PLC intricacies and design details, we resolve the first open problem described in the previous section. We also propose PHY/MAC abstraction models that

accurately capture performance and simplify the complex, exact channel or PHY-layer representations. We compare WiFi and PLC from channel to protocols-implementation perspectives, finding important differences in their spatial and temporal variability and showing that the two mediums complement each other.

Main contribution: We give a crash course on PLC, building a solid background for the remainder of the dissertation and for future research. We focus on factors that affect end-to-end performance and we highlight all fundamental components of PLC from PHY to MAC layer.

Chapter 3 To explore the spatial and temporal variation of single PLC links, we employ a testbed of 19 stations on a realistic enterprise-environment. This study is important for two reasons: (i) future design and deployment of PLC networks and (ii) efficient link-metric estimation in hybrid networks. We show that the variability over space is challenging to predict, due to severe asymmetry in PLC links. Yet, in general, PLC provides very good connectivity, forming almost fully-connected topologies. The variability over time is more stable; we find that most PLC links yield a small capacity deviation and that good links vary less often than bad links. The measurement study of this chapter is further employed for establishing practical link-metric instructions and for facilitating the future deployment of hybrid networks that comprise PLC technologies. We focus on two metrics required by IEEE 1905 standard: the “PHY rate” (capacity) and the “packet errors” (loss rate). We explore link metrics and their variations with respect to space, time, and background traffic.

Main contribution: We present a thorough spatio-temporal study on the PLC capacity variability. Based on this study, we design a capacity-estimation technique and propose systematic link-metric guidelines.

Chapter 4 After studying the performance and variability of single links, we turn our attention to efficiency when multiple users contend for the medium. Analyzing the PLC CSMA/CA protocol is challenging, due to the deferral counter mechanism. Yet, we introduce a model for the multi-user performance of PLC by relying on the so-called decoupling assumption that considers the backoff processes of the stations as independent. The model results in a fixed-point equation with respect to collision probability at the stationary regime of the system. For the first time, we prove that this equation has a unique solution for a wide range of configurations, then we engage our model to propose configurations that enhance performance. We present an extensive performance evaluation with testbed validations of our model,

simulator, and enhancements. We also discuss the convergence time of our enhancement and optimal protocols under dynamic traffic scenarios. Notably, with the proposed configurations, PLC provides a performance very close to that of an optimally configured MAC protocol without knowing the total number of contending stations. In contrast, similar methods for enhancing the WiFi CSMA/CA process do require knowing the number of stations, which complicates their practicality in real deployments.

Main contribution: We introduce a simple model of the PLC CSMA/CA and we employ it to devise configurations that significantly boost throughput for best-effort applications. With our proposed configuration, PLC CSMA/CA represents an interesting step towards a practical protocol that performs close to optimum.

Chapter 5 We delve into the short-term dynamics of the CSMA/CA protocol for PLC. By employing various metrics, we investigate analytically, experimentally and in simulation the short-term fairness of PLC. We prove that PLC is short-term unfair. Interestingly, we find that in PLC CSMA/CA any throughput improvement always comes at a cost of fairness, hence jitter. The key mechanism for this tradeoff is the deferral counter that forces the stations to reduce their access probability when they sense the medium busy. We extensively investigate this compromise, which is crucial for both best-effort and delay-sensitive applications. In addition, we explore the impact of short-term unfairness on delay variance and on the modeling accuracy of the decoupling assumption. To this end, we compare WiFi and PLC protocols. We show that PLC stations are more coupled, compared to WiFi ones, in testbed and simulation; this penalizes the accuracy of our “decoupling-assumption” model for small number of stations. To the best of our knowledge, this is the first experimental evidence for both short-term unfairness and the strong coupling between the PLC stations.

Main contribution: For the first time, we find that the PLC CSMA/CA is by default short-term unfair. We provide fundamental evidence about performance tradeoffs of the PLC CSMA/CA and investigate the effect of each MAC parameter on these tradeoffs.

Chapter 6 After observing that there exists a strong coupling between the stations, which introduces modeling inaccuracies, we employ a more complex multi-user analysis that does not assume independence. We formulate a model for the PLC multi-user performance, called “Drift”, without resorting to the decoupling assumption. The resulting model is a non-linear dynamical system of the expected change in the number of stations at each state

of the system. Compared to the decoupling-assumption model, Drift is more complex, more accurate for small number of stations, and it can capture the transient regime of the system in addition to the stationary one. We provide a wide range of configurations for which the Drift model admits a unique solution, a challenging task due to both protocol and model complexity. Our performance evaluation compares our two models, discusses transient regimes and systems with more than one equilibrium and the accuracy of our model with respect to the stochastic system. Overall, we harness our two models for different applications: (i) We leverage the simple decoupling-assumption model to boost throughput for best-effort and elastic applications (Chapter 4). (ii) We exploit the high accuracy of the Drift model to develop an algorithm that returns a configuration that optimizes fairness and satisfies a throughput constraint, hence tuning the tradeoff that exists in multi-user PLC for delay-sensitive applications.

Main contribution: To the best of our knowledge, the Drift model is the first to reach this level of accuracy and to predict transient performance. We use this accurate model to design an algorithm that exploits the tradeoff between throughput and fairness existing in PLC.

1.4.1 WiFi vs. PLC: Which One Performs Better?

In addition to the contributions described above, the results of this dissertation enable us to compare the performance of wireless and power-line communications. In Figure 1.2, we compare qualitatively the performance of WiFi and PLC, based on our findings.

In terms of time, we find that WiFi has a higher variability compared to PLC. This is due to the intrinsic characteristics of the channels and PHY layers: WiFi is affected by mobility of scattering objects or terminals, whereas PLC is affected by the electrical activity that varies on slower timescales. In case of packet errors, PLC adapts the PHY rate and tackles the frequency selectivity of the channel more efficiently than WiFi. In terms of space, PLC has high spatial-variability due to severe channel asymmetry, whereas WiFi performance can be readily estimated as a function of distance with state-of-the-art path-loss models. In summary, PLC performance is more stable over time, but it is more difficult to estimate which locations yield high capacity, compared to WiFi.

We uncover that the multi-user performance of these two technologies also includes fundamental differences. We find that PLC default configuration is short-term unfair—which introduces high jitter—but more efficient in terms of throughput compared to the one of WiFi. Both short-term and long-term dynamics of the two technologies are diverse, enabling different levels of quality-of-service in a hybrid network of contending flows. Our multi-user study reveals that the PLC CSMA/CA scales better than WiFi with the number of users, which is crucial for capacity planning in dense deployments.

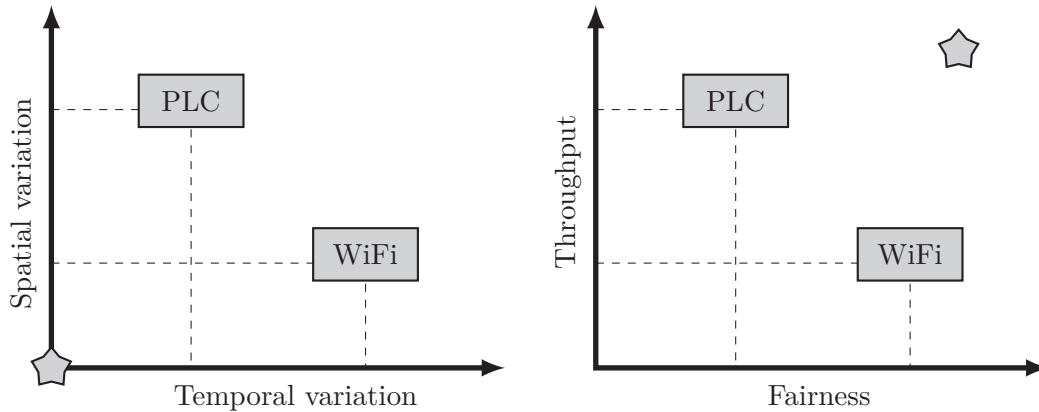


Figure 1.2 – The fundamental observations of this thesis with respect to power-line and wireless communications, presented qualitatively. On the left, the larger the metric is, the worse the capacity stability is. The star represents the ideal, non-varying performance, hence an error-free channel. Spatio-temporal variations are related to the channel and PHY layer; they are discussed in Chapters 2–3. On the right, the larger the metric is, the better the performance is. The star represents a perfect, optimal, and fair resource allocation without collisions. The throughput-fairness tradeoff is related to the MAC layer and access-control protocols that are the subject of Chapters 4–6.

Both single- and multi-user performance findings prove that PLC is a promising technology for residential and enterprise networks, as it complements and very often has better benefits than WiFi. The rich medium-diversity can be beneficial in hybrid networks. In contrast to WiFi, PLC can fully exploit the whole spectrum of the throughput-fairness tradeoff simply by adjusting MAC protocol parameters.

2 A Crash Course on Power-Line Communication

2.1 Introduction

In this chapter, we discuss the main characteristics of PLC, from a channel to protocol perspective. We briefly demonstrate the important components of PLC and enhance the understanding of end-user performance of this unexplored technology. In addition, we build a solid background for the following chapters of this dissertation.

A power-line is a very harsh channel, in terms of noise and attenuation. Compared to the wireless medium, the power-line has diverse types of noise that cannot be modeled as Gaussian. Moreover, there is high frequency-selectivity and attenuation. We review the PLC channel characteristics in Section 2.2. Then, in Section 2.3, we present the most popular specifications and standardizations that succeeded in coping with these harsh channel conditions and in providing promising data-rates.

The peculiar PLC channel properties lead to a more complex design for both PHY and MAC layers compared to the one of WiFi. We give an overview of the PHY layer in Section 2.4. This study is crucial for understanding end-to-end performance and for fully exploiting PLC in hybrid networks later in Chapter 3. We present the MAC layer in Section 2.5, where we focus on the CSMA/CA protocol. This protocol is much more intricate than the one of WiFi, but it has more potential in terms of performance. We extensively investigate it in Chapters 4–6.

Network management and security are very important issues in networking. In Section 2.6, we explore how these problems are engineered in PLC. We unveil the network structure and the management methods. Furthermore, we demonstrate the main aspects of security and station authentication.

In Section 2.7, we compare WiFi and PLC technologies. PLC is inevitably combined with WiFi, because WiFi offers the unbeatable advantage of mobility. In order to fully exploit both technologies, we underline their differences and conduct an experimental study to validate the diversity of the medium and the gains PLC can bring in today's networks, with respect to both coverage and performance variability. We give a summary of this chapter in Section 2.8.

2.2 Channel Characteristics

We explain the main components of the channel, that is, attenuation and noise. They affect both the spatial and temporal variations of PLC capacity. Let us delve into the physical explanations of these impairments. Consider an example of a simple electrical network with a transmitter (TX) and a receiver (RX), as given in Figure 2.1. The main sources of attenuation and noise are the electrical appliances plugged in. Modeled with dashed boxes in Figure 2.1, each connected appliance has an impedance and produces some noise that is non-Gaussian and that depends on the device type [17].




Figure 2.1 – Multi-path and noise in PLC channels.

2.2.1 Attenuation and Frequency Selectivity

The PLC channel is multi-path, because impedance mismatches in the power-line cause multiple reflections of the transmitted signal. These mismatches are due to the electrical appliances attached and depend on the working mode of the appliances, that is, on/off and idle/active. With respect to communications, the electrical cable becomes a transmission line, with a characteristic impedance. The connection of appliances creates impedance mismatches to this transmission line, causing the transmitted signal to be reflected multiple times. For example, in Figure 2.1, at point M, we have an impedance mismatch and any signal s arriving at M is partly reflected (signal r) and partly propagates (signal t) towards the same direction as the original signal s . Reflections of signals at various impedance-mismatched points result in multiple versions of the initially transmitted signal arriving at different times at the receiver, thus establishing a multi-path channel. The delay spread of the PLC channel is typically between 1 and 2 μs , with a few exceptions over 5 μs [18]. The PHY layer of PLC is designed to cope with this delay spread and these multi-path effects, as we observe in Section 2.4.

Multi-path effects cause frequency selectivity and induce more severe attenuation than potential losses due to long cable lengths. The spatial variation of PLC is mainly affected by the position, the impedance, and the number of appliances connected to the electrical grid, rather than by the cable length between the stations—although the number of appliances and the cable length can be correlated. As the aggregated electrical activity varies during the day, the transfer function and attenuation change at large timescales (order of minutes or hours). Hence, temporal variation is affected by multi-path effects, i.e., the appliances' impedances, on long-term timescales, whereas it is affected mainly

by noise on short-term timescales, as we discuss in Section 2.2.2.

Additional Attenuation Factors

In addition to attenuation introduced by electrical appliances, there are additional attenuation factors that are usually not considered in channel modeling and that arise due to the structure of the electrical grid and power distribution. For instance, significant attenuation is introduced between a pair of stations connected to different circuit breakers, phases or distribution boards. The impedance of the appliances and these extra attenuation factors render the prediction of PLC link-quality challenging, given the location of two stations. In Section 3.2, we discuss these factors and we observe in our testbed that spatial-variation of PLC is challenging to predict.

2.2.2 Types of Noise

There are various types of noise generated in a PLC channel. A widely accepted classification of noise (e.g., [17, 19]) identifies the following types:

Colored background noise is due to different low-power noise sources present in the network and it is usually characterized with a power/spectral density decreasing with frequency.

Narrowband interferences are caused by radio broadcasters and are usually sinusoidal signals.

Impulsive noise is the most severe and intricate type of noise, because each electrical appliance exhibits unique characteristics, from both time and frequency perspectives. It is further distinguished into three classes that we describe next.

The main sources of impulsive noise in PLC are electrical appliances: they create synchronous or asynchronous noise to AC mains and aperiodic noise when switched on/off. In detail, the three classes of impulsive noises are as follows.

Periodic, synchronous to the mains noise is a cyclostationary noise, synchronous with the mains and with a frequency of 50/60 Hz (depending on the country). It is often generated by rectifier diodes used in electrical appliances.

Periodic, asynchronous to the mains noise is generated mostly by switched power supplies and has a frequency of tens or hundreds kHz (in the range 12–217 kHz [19]).

Aperiodic noise is caused by people switching on and off electrical appliances and plugging them in.

The short-term variation of the channel is due to noise. Two timescales describing this variation, called *invariance* and *cycle* scales, are introduced by Sancha et al. [20]. An invariance scale is defined at a level of subintervals of the mains cycle and is related to the noise synchronous to mains, whereas the cycle scale represents different mains cycles and is affected by noise asynchronous to mains. PLC modulation and coding are tailored to the noise variation within or across the mains cycle, as we explain in Section 2.4.

2.3 Standardization and Specifications

We present the most popular PLC specifications that tackle the challenging channel conditions described in the previous section. The HomePlug Alliance has released various specifications, from 2001 until today. Table 2.1 shows all specifications and IEEE standardization efforts. The first specification, 1.0.1, supports rates up to 14 Mbps. The most recent one, AV2, reaches 1.5 Gbps rates; it was released in 2012, revealing the rapid development and evolution of PLC.

Specification/Standard	Rate supported	Purpose
HomePlug 1.0.1	14 Mbps	“No-new-wires” connectivity
HomePlug Turbo	85 Mbps	Standard definition video, IPTV, multi-room DVR network, etc.
HomePlug AV	200 Mbps	Applications such as HDTV and VoIP
HomePlug AV500	500 Mbps	Enhanced AV rates, by bandwidth extension
IEEE 1901	500 Mbps	Standardization
HomePlug Green PHY	10 Mbps	Automation, comfort and smart grid
HomePlug AV2	1.5 Gbps	4K Ultra HD video and bandwidth-hungry applications

Table 2.1 – PLC specifications and standardization in ascending date-order.

In 2010, IEEE published the 1901 standard [10] for broadband PLC in the band of 1–30 MHz, with a possibility of an extension to 50 MHz. Compared to the HomePlug specifications, IEEE 1901 standardizes both in-home and outdoor-access PLC networks. The relationship between HomePlug and IEEE 1901 is similar to that between Wi-Fi Alliance and IEEE 802.11 standard [21]. HomePlug certifies PLC products to conform to the 1901 standard, and enhances specifications at higher rates that are fully-interoperable with previous solutions.

In addition to solutions for bandwidth-hungry applications, HomePlug released the Green PHY specification for home automation, electric vehicle communication, and the smart grid. Green PHY operates at the same frequency band as the other HomePlug specifications, hence it is fully interoperable and can successfully interact with them. Compared to high-rate solutions, Green PHY typically offers larger coverage, lower rates and lower power-consumption.

In the following sections, we review the main features of the PHY and MAC layers for the most popular PLC specifications, HomePlug AV and AV2.

2.4 PHY Layer

In this section, we discuss the main PHY features of broadband PLC and we distinguish the differences between the HomePlug AV and AV2 specifications. We first describe the modulation and coding schemes. Then, we uncover how these schemes are chosen per link, that is, the channel estimation and adaptation process.

2.4.1 Modulation and Coding

To cope with the multi-path nature of PLC, the PHY layer of HomePlug AV is based on an FFT-OFDM scheme with 1155 carriers in the 1.8–30 MHz frequency band. The whole band is used for every frame transmission, with only a few carriers notched out to avoid interference to other existing technologies, such as amateur radio bands. Each OFDM carrier can employ a different modulation scheme among BPSK, QPSK, and 8/16/64/256/1024-QAM, because of the frequency selectivity of the channel. The OFDM symbol duration is 40.96 μ s and, to overcome the delay spread of the channel, there is a variable guard interval that can take the values 5.56, 7.56 or 47.12 μ s. The forward error correction code (FEC) used is turbo convolutional coding (TCC) and the available rates are 1/2 and 16/21¹. TCC is succeeded by channel interleaving for impulsive noise, prevalent in PLC channels. The combination of OFDM, guard interval with cyclic prefix, TCC and interleaving enable robust and high-rate communications over a challenging channel.

HomePlug AV2: New Key Features

PLC modulation and coding techniques are continuously evolving. The most recent specification, AV2, introduces additional features to support bandwidth-hungry applications and improved data-rates. Compared to AV, it incorporates the following:

- extension of the frequency spectrum to 1.8–86 MHz,
- 4096-QAM modulation scheme, 8/9 FEC rate, and smaller guard intervals,
- lower MAC overhead via shorter delimiters and delayed acknowledgments,
- support for MIMO (multiple-input and multiple-output) with beamforming.

The most significant AV2 feature is MIMO transmission. Whereas HomePlug AV uses the line-neutral cable pair, AV2 employs any two pairs formed by the line (L), neutral (N) or protective-earth (PE) wires (i.e., L-N, L-PE or N-PE). This allows for significantly improved peak data-rates and performance. MIMO uses two independent transmitters and up to four receivers, with beamforming required to maximize the performance on the independent streams. A receiver plugged in to power outlets that are connected to the three-wire line, neutral, and protective-earth can employ up to three differential mode receive ports and one common-mode port. The common-mode signal is the voltage

¹To achieve robust, low-rate communication, HomePlug Green PHY (the home-automation specification introduced in the previous section) uses the same frequency band as AV and only the QPSK scheme and the 1/2 rate.

difference between the sum of the three wires and the ground. PLC MIMO requires that the third wire, i.e., PE, be installed at the building, which is not the case for a few old electrical installations. Schwager presents a detailed survey on AV2 features and MIMO [22].

Physical Blocks

Before modulation and transmission, the data in PLC is segmented into 512-byte blocks, called physical blocks (PBs)². Each PB has a 8-byte header for identification and cyclic-redundancy-check purposes. PBs are individually encoded with the TCC and typically a large number of PBs are aggregated into a MAC protocol data unit (MPDU). By efficient retransmission of only the corrupted PBs with selective acknowledgments and an automatic-repeat-request (ARQ) protocol, impulsive noise in PLC is handled. The standard [10] enables PBs to be retransmitted multiple times for robustness (even within a single MPDU), but the detailed ARQ protocol is vendor-specific. PBs are the fundamental units of PLC transmission and will be very important in estimating error rates and the actual capacity.

2.4.2 Channel Estimation and Adaptation

Because each carrier can employ a different modulation scheme, PLC stations exchange messages with the scheme that each carrier uses, the forward error correction code (FEC) rate, and other PHY-layer parameters. The entity that defines these PHY options is called the tone map, and it is estimated during the channel-estimation process.

Channel estimation is performed if and only if the station has data to transmit. The source initially sends “sound” frames to the destination by using a default, robust modulation scheme that employs QPSK for all carriers. QPSK is used for the initial channel estimation and communication between two stations, but also for all broadcast and multicast transmissions. After estimating the tone map, the receiver sends the tone map with a unique identification back to the source, called the tone-map index. The source uses the tone-map index and tags each frame with it to inform the receiver on the demodulation procedure. The receiver can choose up to $L + 1$ tone maps, for some integer L : L tone maps for different sub-intervals of the AC-line cycle, called slots, and one default tone map. PLC uses multiple tone-maps for the different slots of the AC-line cycle to cope with the periodic impulsive noise that is synchronous to the mains.

Katar et al. [23] investigate the optimal duration of the tone-map slots, and they conclude that 1–2 ms duration can yield significant gains. In this dissertation, our experimental results unveil that, in practice, there are $L = 6$ tone maps chosen and used in the first half of the AC line cycle and then repeated periodically, i.e., with period 10 ms for 50-Hz mains. This yields that, indeed, the tone-map slot has a duration of 1.67 ms.

²In addition, 136-byte and 16-byte block sizes are supported.

Tone maps are updated dynamically, either when they expire (after 30 s) or when the error rate or SINR exceed a threshold. When channel conditions change, the receiver adapts the modulation scheme per carrier and other PHY layer parameters. The channel-estimation algorithm is not defined by the IEEE 1901 standard, but it is vendor-specific. Due to the proprietary nature of these algorithms, we have not yet managed to gain access to them.

The Bit-Loading Estimate

The bit-loading estimate (BLE) is an estimation of the number of bits that can be carried on the channel per μs . BLE is retrieved from the frame control, called the start-of-frame (SoF) delimiter. Each MPDU is preceded by its SoF³. The SoF contains information for both PHY and MAC layers, such as BLE or MPDU duration. BLE is defined as follows.

Definition 1. [10] *Let T_{sym} be the OFDM symbol length in μs (including the guard interval), R be the FEC code rate, and PB_{err} be the PB error rate (chosen based on the expected PB error rate on the link when a new tone map is generated. It remains fixed until the tone map becomes invalidated by a newer tone-map). Let also B represent the sum of number of bits per symbol over all carriers. Then, BLE is given by*

$$BLE = \frac{B \times R \times (1 - PB_{err})}{T_{sym}}. \quad (2.1)$$

BLE takes into account the overhead due to FEC but does not include any MAC-layer overhead, such as frame-control duration or inter-frame spaces. It is an estimation of the capacity, as we observe in Section 3.4. The bit loading is unique per pair of stations and it can vary within the mains cycle, as we explained above.

2.5 MAC Layer

We now review the PLC MAC layer and its most important sub-functions. PLC enables the aggregation of multiple Ethernet packets into one PLC frame. We first describe the frame aggregation process from a MAC-layer perspective and provide some level of abstraction for facilitating the modeling and simulation of this complex procedure. Then, we discuss in detail the CSMA/CA protocol of PLC, which is the main access control mechanism. This protocol is defined in the IEEE 1901 standard, hence we refer to it as “IEEE 1901”.

³Before transmission, the MPDU is converted to a physical protocol data unit (PPDU), which is the actual entity transmitted over the power-line. The PPDU adds a preamble before the SoF in order to ensure detection by other stations and backward compatibility with previous PLC specifications.




Figure 2.2 – The frame aggregation and acknowledgment process in PLC.

2.5.1 Frame Structure and Aggregation

Frame structure and aggregation are important for the end-to-end performance of PLC. To reduce various types of overhead, multiple Ethernet packets are transmitted with a single PLC frame. We answer two questions here: (i) How many Ethernet packets are contained in a PLC frame? (ii) How can the end-to-end performance be accurately simulated without the enormous complexity of channel modeling and PHY layer?

From a performance or simulation perspective, the metrics of interest of the frame-aggregation process would be the number of Ethernet packets per MPDU and the MPDU duration. We now describe the MAC layer processes that determine these metrics and that are sketched in Figure 2.2. As we mentioned in Section 2.4, the Ethernet packets are organized in PBs. Then, the PBs are forwarded to a queue, and based on the BLE of the current tone-map slot s , BLE_s , they are aggregated into an MPDU. The MPDU duration is determined by BLE_s , the maximum MPDU duration (specified by [10]), and an aggregation timer that fires every few hundreds of ms after the arrival of the first PB, as concluded from our testbed measurements⁴. The modulated MPDU is transmitted by a CSMA/CA protocol described in the next section. The receiver demodulates and decodes the MPDU and transmits a SACK that informs the transmitter about which PBs were received with errors. In Chapter 3, we observe that the full retransmission and aggregation process, and, as a result, the MAC and PHY layers, can be modeled using only two metrics that are defined in Definition 1, that is, PB_{err} and BLE_s . The testbed measurements of these metrics can be plugged into any simulator and an assumption on the error distribution would be required to model the PB errors. Thus, an exact channel representation per link and a PHY layer implementation can be circumvented.

In the case where an MPDU reaches the maximum frame duration and there is high traffic intensity, multiple MPDUs can be transmitted. In particular, stations are

⁴Note that the MPDU duration is a multiple of the OFDM symbol length, and that padding is used to fill these symbols.

allowed to transmit multiple MPDUs in a burst via CSMA/CA access. Bursts contend for the medium and not individual MPDUs, in this case. The number of MPDUs per burst depends on channel conditions and station capabilities. Up to four MPDUs may be supported in a burst. A selective acknowledgment (SACK) of the burst acknowledges each PB, so that only the corrupted PBs are retransmitted. If a frame is received with all the PBs corrupted, then it is considered to have collided with another frame.

With the above simplification of the frame aggregation process, PLC can be evaluated by using PB_{err} and BLE_s , and the CSMA/CA protocol that we describe next. This is very important for simulating hybrid networks, as the exact channel-model and PHY-layer representations are quite challenging with PLC.

2.5.2 Access Control: The IEEE 1901 CSMA/CA Protocol

The MAC layer of IEEE 1901 includes both TDMA and CSMA/CA protocols [10]. However, to the best of our knowledge, all current commercial devices implement only CSMA/CA. The same CSMA/CA protocol is used for all specifications described in Table 2.1, hence it is prevalent in PLC networks. This CSMA/CA protocol is similar to the extensively studied 802.11 for wireless communications, but with important differences. The main difference is that, contrary to WiFi, PLC stations increase their contention windows (CW) not only after a collision, but also after sensing the medium busy. This is regulated by an additional counter, called the deferral counter. In this section, we present the relevant aspects of the 1901 CSMA/CA procedure. By comparing 1901 with 802.11, we gain insights on the requirements that established this intricate protocol.

The first HomePlug specification that included this CSMA/CA mechanism is HomePlug 1.0. HomePlug 1.0 employs a frame preamble that comprises 7.5 OFDM symbols. The slot duration was determined by the time required by a station to decide whether the medium is busy or idle (i.e., to detect a preamble transmission), and it is equal to the duration of 7 HomePlug 1.0 symbols (i.e., $35.84 \mu s$). Although newer technologies have different symbol durations, the slot duration has remained the same for all HomePlug standards for backward compatibility. Observe that the slot duration is large compared to the one of 802.11 (which is $9 \mu s$ for 802.11a/g/n/ac). In the next paragraphs, we explain the effect of the slot duration in the backoff process.

The backoff process of 1901 uses two counters: the backoff counter (BC) and the deferral counter (DC). In addition, there are four *backoff stages*⁵. Now, we discuss the common features of 1901 and 802.11, and later we elaborate on the deferral counter. When a new packet arrives for transmission, the station starts at backoff stage 0, and it draws the backoff counter BC uniformly at random in $\{0, \dots, CW_0 - 1\}$, where CW_0 refers to the contention window used at backoff stage 0. Similarly to 802.11, BC is decreased by 1 at each time slot if the station senses the medium to be idle, and it is

⁵In the standard [10], the backoff stage is determined by the so-called *backoff procedure counter* (BPC).

Priority class:	CA0/CA1		CA2/CA3	
backoff stage i	CW_i	d_i	CW_i	d_i
0	8	0	8	0
1	16	1	16	1
2	32	3	16	3
3	64	15	32	15

Table 2.2 – IEEE 1901 parameters for the contention windows CW_i and the initial values d_i of deferral counter DC , for each backoff stage i .

frozen when the medium is sensed busy. In the case where the medium is sensed busy, BC is also decreased by 1, once the medium is sensed idle again. When BC reaches 0, the station attempts to transmit the packet. Also similarly to 802.11, the station jumps to the next backoff stage if the transmission fails (unless it is already at the last backoff stage, in which case it re-enters this backoff stage). When entering backoff stage i , a station draws BC uniformly at random in $\{0, \dots, CW_i - 1\}$, where CW_i is the contention window at backoff stage i , and the process is repeated. For 802.11, the contention window is doubled between two successive backoff stages, thus $CW_i = 2^i CW_0$. For 1901, CW_i depends on the priority level and is given in Table 2.2. There are four priority classes in 1901, CA0 to CA3. CA0/CA1 priorities serve best-effort applications, and CA2/CA3 serve the delay-sensitive ones.

Now, when there are few contending stations (i.e., 1 or 2), or when the traffic load is very low, the time spent in backoff causes a large overhead and increases as the contention window increases. Given the large slot-duration of 1901, the average delay due to backoff ($\sim (CW_0 - 1)/2$ time slots) can be reduced when there are few contending stations—that is, small collision likelihood—by choosing a small minimum contention-window, e.g., $CW_0 = 8$, as specified for 1901 (Table 2.2). However, as expected, small contention-windows yield higher collision probabilities when the number of stations increases or when the traffic load rises. The deferral counter DC was introduced as a countermeasure in the CSMA/CA process of 1901, to reduce collisions induced by small contention-windows. This is achieved by triggering a redraw of the backoff counter BC *before* the station attempts a transmission.

The main difference between 1901 and 802.11 is the introduction of DC in 1901, which enables a station to enter a higher backoff stage even if it did not attempt a transmission. The mechanism for deciding when this occurs works as follows. When entering backoff stage i , DC is set at *an initial DC value* d_i , where d_i is given in Table 2.2 for each i . After having sensed the medium busy, a station decreases DC by 1 (in addition to BC). If the medium is sensed busy and $DC = 0$, then the station jumps to the next backoff stage (or re-enters the last backoff stage, if it is already at this stage), and it re-draws BC *without attempting a transmission*. Figure 2.3 shows an example of such a backoff process.

		Station A			Station B				
		backoff stage i	CW_i	DC	BC	backoff stage i	CW_i	DC	BC
time ↓	$i = 0$		8	0	3	$i = 0$	8	0	5
			⋮	⋮	⋮			⋮	⋮
			8	0	0		8	0	2
			Transmission						
			8	0	7	$i = 1$	16	1	11
			⋮	⋮	⋮			⋮	⋮
		8	0	0			16	1	4
		Transmission							
		8	0	5		16	0	3	
		⋮	⋮	⋮		⋮	⋮	⋮	
	8	0	2		16	0	0		
	Transmission								
	$i = 1$	16	1	6	$i = 0$	8	0	2	
	⋮	⋮	⋮			⋮	⋮	⋮	

Figure 2.3 – Evolution of the IEEE 1901 backoff process with 2 saturated stations A and B . Initially, both stations are at backoff stage 0. A transmits twice consecutively. Note the change in i when a station senses the medium busy and has $DC = 0$.

We now discuss the effect of the deferral counter on jitter, because this justifies the choice of rest of the parameters in 1901. As we have seen above, the deferral counter manages to reduce collisions. It can however introduce high jitter. The transmitting station has an advantage over the other stations because its CW is maintained at its initial value CW_0 , whereas the other stations might increase their CW when $DC = 0$, as explained earlier. As a result, a station can grasp the channel for multiple consecutive transmissions, until it releases the channel; and then it can wait for a long sequence of other station transmissions, until it obtains the token again, which yields high jitter. To reduce the CW imbalance between the transmitting station and the others, 1901 employs a total of four backoff stages. Moreover, there are two different configurations for best-effort and delay-sensitive applications (see Table 2.2). The delay-sensitive class employs smaller contention-windows, and the contention window is not doubled between backoff stages 1 and 2. This improves jitter, but yields a higher collision probability, i.e., lower throughput, compared to the CA0/CA1 classes. Figure 2.4 summarizes the design choices concerning the MAC parameters described in this section. We focus on best-effort applications hence, throughput in Chapters 4 and 6 and on delay-sensitive applications in Chapters 5 and 6. There, we investigate in detail the effect of the MAC parameters on throughput and fairness (or jitter).

Quality-of-Service Differentiation and Priority-Resolution Process

As we observe above, PLC MAC layer uses four priority classes for best-effort and delay-sensitive applications, which are CA0 to CA3, from lowest to highest priority. The CA0/CA1 classes have different MAC parameters than CA2/CA3 classes. In addition,

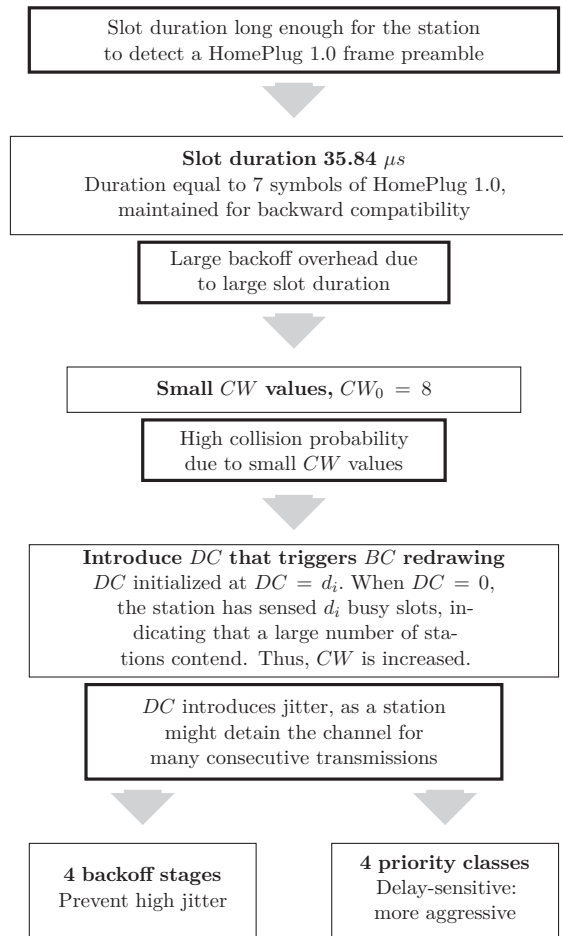


Figure 2.4 – A discussion of the design choices of the IEEE 1901 CSMA/CA.

the standard specifies that there is a timeout on the frame transmission that is vendor-specific, and the default value in our hardware was found to be 2.5 s for the CA0/CA1 class and 0.3 s for the CA2/CA3. The packet priority depends on the application. In Appendix B, we explain how to modify the PLC priority of a flow.

The standard specifies that only the stations belonging to the highest priority can contend for the medium through the CSMA/CA procedure introduced in this section. The highest priority is determined using two priority-resolution slots, called PRS0 and PRS1, in which a station with packets of a specific class either transmits or senses the medium for a higher priority signal. CA3, the highest priority stations, transmit on both PRS0 and PRS1. CA2-stations signal only on PRS0, and CA1-stations transmit only on PRS1 and on the condition that PRS0 is empty. In this way, the rest of the stations sense these transmissions and infer whether a station with higher priority aims to transmit. In which case, they defer their transmission. The priority-resolution slots follow every successful transmission in the channel.

2.6 Network Structure and Management

We explain the PLC network structure and the crucial features of management.

Central Coordinator

PLC uses a centralized authority, called the central coordinator (CCo), to manage the network. To operate, each station must join a network with a CCo; however, transmissions do not need to go through CCo. Usually, CCo becomes the first station that is plugged and it can change dynamically, if another station has better channel capabilities than current CCo does. CCo is responsible for the authentication and association of new stations, and for allocating TDMA and CSMA/CA sessions. It also transmits beacons periodically so that the stations are synchronized with the AC line cycle. The beacons are vital for both PHY and MAC layers, as explained in the next paragraph.

The CCo is chosen and, in case of failure, updated automatically. Thus, the user does not have to set a CCo or to handle CCo changes. HomePlug aims to provide seamless network operation under CCo modifications and to hide the CCo intricacies from the user. All stations behave the same from an end-user perspective, whether they act as a CCo or not. In Appendix B, we give guidelines on how to set the CCo statically, in case where a steady network-structure is desired, such as enterprise environments.

Beacons

The CCo of the PLC network transmits beacons to synchronize the stations with respect to the mains cycle and to manage the network. The beacon interval is twice the period of the mains cycle. Usually, the beacon is transmitted within a beacon region that follows every second positive zero-crossing of the AC line cycle. The beacon interval can accommodate several beacons in case multiple PLC networks are within sensing range of each other, and it enables multiple beacons to be transmitted without collision.

The beacon is overheard by the stations and is used for various purposes. First, the stations synchronize their periodic modulation-scheme (tone maps) over the mains by using the beacon. Second, the stations get notified about TDMA and CSMA/CA allocations within the beacon period. Finally, a few management details and vendor-specific information are transmitted within the beacons.

Security and Authentication

Power-line signals can propagate from one cable to another, thus privacy and security issues are raised with PLC. To address these issues, security mechanisms built on 128-bit advanced encryption standard (AES) are employed in PLC. The stations have to be configured with a network membership key (NMK), that acts similarly to the password for security in WiFi. Only stations that share a common NMK can exchange data.

The CCo ensures security and privacy in the network, by associating and authenticating stations. First, during the association procedure a station obtains a terminal equipment identifier (TEI) from the CCo. It then provides the station with a network encryption key (NEK), as long as the station has a valid NMK, which is the authentication phase. The station encrypts each PB with the NEK. We describe how to set the NMK of a station, in Appendix B. Latchman et al. present a detailed description of association and authentication [18].

Management Messages

Management messages (MMs) are a key feature and tool for PLC. They are used for network management, authentication, association, tone-map establishment and updating. Stations must exchange MMs each time the tone map is updated, because the source has to be notified for the modulation scheme that each carrier uses and other PHY parameters. If MMs are transmitted through the PLC channel, they are tagged either with CA2 or CA3 priority, because they are crucial for the network operation.

In addition to being transmitted between stations, MMs can be transmitted from the Ethernet interface to the PLC chip. This enables us to interact with the PLC stations and to configure the devices or measure statistics, as we explain in Appendix B.

2.7 WiFi vs. PLC

We summarize the main differences between WiFi and PLC technologies. These differences span the channel nature and the PHY, MAC layers. Table 2.3 compares the two technologies. As we observe, WiFi and PLC use different frequency spectra, with WiFi employing non-interfering channels for transmission with 20, 40, 80 or 160 MHz bandwidth—depending on the 802.11 version—and PLC using the entire available spectrum. Both mediums are multi-path, that leads to an OFDM-based PHY layer with PLC that allows each carrier to employ a different modulation scheme in order to tackle high frequency-selectivity. Due to this feature, PLC rate-selection is performed at the receiver that has SINR samples over the whole spectrum. In contrast, in state-of-the-art WiFi solutions, the transmitter estimates the optimal data-rate, based on the acknowledgments it receives and a moving average of the packet error rate. These differences can lead to potential high overhead for rate-selection in PLC, because this procedure necessitates message-passing per station-pair. However, compared to WiFi, it can yield more efficient adaptation to channel errors, as we will notice in our experimental study in the next paragraph.

The frequency and time selectivity are affected by different factors for the two technologies. This diversity can be beneficial for augmenting network reliability in residential and enterprise settings. In the following sections, we conduct an experimental study to explore to which extent we can augment WiFi network performance, by introducing PLC. This validates in practice the rich diversity between the two technologies.

Feature	WiFi	PLC
Frequency bands	2.4 or 5 GHz	1–86 MHz
Channelization	Yes	No
Channel nature	Multipath	Multipath
Frequency-selectivity causes	Reflection, diffraction, scattering of signal	Electrical appliances impedance, circuit breakers, phase coupling
Time-selectivity causes	Mobility of terminals or scattering objects, fading	Impedance variation and noise introduced by electrical appliances
OFDM	All carriers use the same scheme	Each carrier may use a different scheme
Rate selection	At transmitter, by using packet error rate (PER)	At receiver, with feedback to the transmitter
Network structure	Infrastructure (transmissions must pass through access point) or ad-hoc mode	Central coordinator (CCo) obligatory, but transmissions do not need to go through CCo
Medium nature	Shared (broadcast)	Shared (broadcast)
CSMA/CA	Doubles CW only after collision	Doubles CW after collision and possibly also after a sensed transmission

Table 2.3 – WiFi vs. PLC: Comparison of the main characteristics and features.

As we observe in Section 2.5.2, the CSMA/CA mechanisms of the two technologies are different. The CSMA/CA of PLC is the subject of Chapters 4–6, where we compare it with the one of WiFi from various aspects, such as throughput and fairness. Here, we focus only on the temporal and spatial diversity of the two technologies.

2.7.1 Spatial Variation

We first compare the spatial variation of WiFi and PLC in our testbed, where WiFi and PLC interfaces have similar nominal capacities⁶. Our testbed, consisting of 19 stations, is described in Appendix B. This study quantifies the gains that PLC can yield in situations with wireless “blind spots” or bad links and examines which medium an application should use. We conduct the following experiment: For each pair of stations, we measure the available throughput of both mediums back-to-back for 5 minutes, at 100-ms intervals. These experiments are carried out during working hours to emulate a realistic residential/enterprise environment. We show the average and standard deviation

⁶We use 802.11n, with 2 spatial streams, 20 MHz bandwidth and 400 ns guard interval, yielding a maximum PHY rate of 130 Mbps. We selected a frequency that does not interfere with other wireless networks in our building. For PLC, we use HomePlug AV. The highest PLC data-rate is 150 Mbps hence, both interfaces have similar nominal capacities. This is confirmed by the maximum throughputs exhibited by both mediums, shown in Figure 2.5.




Figure 2.5 – WiFi vs. PLC performance for all links (top). Spatial variation of the performance ratio between WiFi and PLC (bottom).

of these measurements (for links with a non-zero throughput for at least one medium).

Let T_W and σ_W be, respectively, the average value and standard deviation of throughput for WiFi (T_P and σ_P , respectively, for PLC). Figure 2.5 illustrates the results of our experiment. Our key findings are as follows.

Connectivity: PLC yields better connectivity than WiFi. 100% of the station pairs that are connected with WiFi are also connected with PLC. In contrast, 81% of the station pairs that are connected by PLC links, are also connected by WiFi links. At long distance (more than 35 m), there is no wireless connectivity, whereas PLC offers up to 41 Mbps. Thus, PLC can eliminate, to a large extent, blind spots in WiFi coverage.

Average performance: 52% of the station pairs exhibit throughput higher with PLC than with WiFi; PLC can achieve throughput up to 18 times higher than WiFi (40.1 vs. 2.2 Mbps). The maximum gain of WiFi vs. PLC was similar, i.e., 12 times (46.3 vs. 3.8 Mbps). Therefore, the two mediums are equivalent and they complement each other, yielding augmented reliability in hybrid settings.

Variability: At short distances (less than 15 m), WiFi usually yields higher through-




Figure 2.6 – Temporal variation of capacity for PLC and WiFi for two links during working hours (started time written).

put, but PLC offers significantly lower variance. WiFi has higher variability with the maximum standard deviation of throughput being $\sigma_W = 19.2$ Mbps vs. $\sigma_P = 3.8$ Mbps for PLC. The vast majority of PLC links yield a σ_P smaller than 4 Mbps. These results are due to the medium nature of both technologies: At short distances, the path loss of WiFi is low, whereas PLC performance strongly depends on the number and type of electrical appliances attached between the stations and on the stations’ circuit breakers or phase.

Conclusion: At long distances, PLC eliminates wireless blind spots or bad links, yielding notable gains. At short distances, although WiFi provides higher throughput, PLC provides significantly lower variance, which can be beneficial for TCP or applications with demanding, constant-rate requirements, such as high-definition streaming. We explain this difference by the ability of PLC to adapt each carrier to a different modulation scheme, contrary to WiFi (see Table 2.3). PLC reacts more efficiently to frequency-selective and bursty errors than WiFi that has to lower the rate for all carriers. In addition, WiFi is affected by the mobility of scattering objects (e.g., mobility of people), whereas PLC channel quality significantly changes when electrical appliances are switched on or plugged in, events that generally occur with lower frequency compared to human mobility.

2.7.2 Temporal Variation

We now look at the concurrent temporal variation of WiFi and PLC during working hours for a much longer duration than before. We are interested in exploring the timescales at which the two mediums vary. Figure 2.6 shows the capacity for concurrent tests on WiFi and PLC, averaged over 50 packets. We observe that link 3-8, which is a good-quality link, exhibits a variation much higher with WiFi than with PLC. Although we would expect

channel changes due to switching electrical appliances in the building, the PLC link is almost not affected by people leaving the premises (around 6pm). The average-quality link 3-0 varies more for both mediums.

These preliminary results indicate that PLC has low variability for good links and high variability for bad links. We observe that, in addition to improving average performance and aggregating bandwidth, we can simultaneously exploit both mediums over short-scales to provide reduced aggregated variability and augmented reliability for demanding applications. To this end, a thorough evaluation of PLC is required, as WiFi is extensively studied. In the following chapter, we focus on PLC technology and explore its temporal and spatial variation.

2.8 Summary

In this chapter, we provide a tutorial on PLC. First, we identify the main features of the channel in order to understand later the end-to-end user performance. Second, we point out the most popular standardizations of commercial solutions to date. These solutions cope with challenging channel conditions and provide high data-rates and robust communication.

We delve into the PHY and MAC layers of PLC, constructing a solid background for the rest of this dissertation. Notably, we observe that PLC design is much more complex than that of WiFi and we justify PLC intricacies. We uncover the most important differences between the two technologies and validate our study by testbed experiments. Our results confirm the potential of PLC to significantly augment network reliability and to extend coverage.

3 Spatio-temporal Variations and Estimation of Capacity

3.1 Introduction

In this chapter, we thoroughly evaluate spatio-temporal variations of end-to-end PLC performance of *single* links, and we propose practical link-metric estimation methods. We delve into both PHY and MAC layers of PLC, via a testbed of more than 140 links, and we verify our findings by using devices from two different vendors and HomePlug specifications. This study is important from both end-user and hybrid-network perspectives: It eases future PLC deployment and the incorporation of PLC into hybrid networks.

In Section 3.2, we uncover the network topology, discuss connectivity, and investigate the spatial variation of PLC. Notably, we find that PLC links are highly asymmetric and this has two consequences: (i) Link metrics should be carefully estimated in both directions. (ii) Predicting which PLC links will be good is challenging. The spatial-variation study unveils important observations for large-scale deployment and connectivity.

We explore the temporal variation of the PLC channel in Section 3.3, and distinguish three different timescales for capacity variability, called “invariance”, “cycle” and “random”. Each timescale yields fundamental insights for exploiting PLC to its fullest extent, for predicting end-user performance variability, and for efficiently updating link metrics (e.g., high-frequency probing yields accurate estimations, but high overhead). We discover that good links vary much less often than bad ones, and that overhead can be reduced by probing good links less often.

In Section 3.4, we introduce a capacity-estimation technique for PLC links and evaluate it under various settings; our capacity estimator is the bit-loading estimate *BLE* (introduced in Section 2.4.2) and included in every PLC frame header. We examine the ideal packet-size and priority-class of probing traffic. We also validate the accuracy of our capacity-estimation technique by designing a load-balancing algorithm for hybrid WiFi/PLC single-hop networks.

To explore metrics related to packet errors and delay, in Section 3.5 we examine the retransmission procedure and how link metrics are affected by contention. There, we investigate both broadcast and unicast methods for evaluating link quality and we

introduce a technique to address potential link-metric sensitivity to background traffic. We also measure PB_{err} that is defined in the IEEE 1901 standard.

Both our metrics BLE and PB_{err} are compliant with IEEE 1905 standard for hybrid networks. We summarize our guidelines for link-metrics estimation in Section 3.6. We review the related work in Section 3.7 and give concluding remarks in Section 3.8.

3.2 Spatial Variation of PLC

The spatial variation of PLC is interesting for future deployment and popularity of PLC and implementation of hybrid networks. We first discuss PLC topology and connectivity, then the predictability of link quality with respect to electrical-cable length.

3.2.1 Topology and Connectivity

In this chapter, our experimental results are conducted on a testbed spread over the second floor of BC building at EPFL. The map of our testbed and the electrical plan of the floor is shown in Figure 3.1. We observe that there are two electrical distribution boards denoted by B1 and B2. Stations 0–11 are connected to B1, and 12–18 are connected to B2. We use two types of hardware with HomePlug AV and AV500 specifications. Compared to AV described in Section 2.4, AV500 extends the bandwidth to 1.8–68 MHz and maintains the rest of the AV PHY-layer mechanisms unchanged. The details on configuring PLC devices and measuring statistics can be found in Appendix B. In the following, we present our findings, with respect to the topology and connectivity of PLC.

As mentioned in Section 2.6, a PLC network requires a central authority called the CCo. Our floor has two distribution lines that are connected with each other at the basement of the building. This means that the cable distance between the two boards (more than 200 m) makes the PLC communication between two stations at different boards challenging. Due to the two distribution lines, none of the stations can communicate with all stations and become the CCo. Hence, we create two PLC networks, represented with different colors in Figure 3.1, with CCo's at stations 1 and 15. This means that in large-scale deployments, a CCo per distribution board must be provisioned due to the potential attenuation introduced between different distribution boards.

In terms of connectivity and topology, we observe that each of the two PLC networks forms an almost fully-connected graph. This yields very good connectivity for PLC networks that belong to a single distribution board. In addition, we observe that there is some connectivity between stations at different boards, e.g., station 11 with 13 and 12. We did not manage to verify whether this connection is due to a direct path between cables (who meet at the basement of our building where two boards are connected) or to electromagnetic interference from neighboring electrical cables on the testbed floor. For this reason, in the following section, we investigate link quality as a function of cable distance, excluding these few cases of links between different distribution boards.




Figure 3.1 – The electrical plan and the stations (0–18) of our testbed. Stations marked with the same color belong to the same network and are connected to the same distribution board (either B1 or B2).

3.2.2 Asymmetry in PLC Links and Consequences

We explore the spatial variation of PLC, as it is important for predicting coverage and the good locations for PLC stations, and for implementing link metrics. We find that PLC is highly asymmetric, and this should be considered when estimating link metrics.

A very important characteristic of power-line channels is that they exhibit performance asymmetry, i.e., capacity can differ significantly between the two directions of the link. In all the experiments we run (both with AV and AV500), we observe a performance asymmetry of more than 1.5x in approximately 30% of stations pairs in our testbed. Figure 3.2 presents typical examples of these links, for which the throughput in one direction is less than 60% of the throughput in the opposite direction. By re-conducting the experiments with AV500 devices, we verify that the asymmetry is not due to the hardware. Link asymmetry in PLC has been also observed by Murty et al. [24]. We attribute this asymmetry to a high electrical-load (for instance, one or more appliances with much higher impedance than the cable’s impedance) that exists near one of the two stations. In this case, the channel cannot be considered as symmetric and the two transmission directions in the link experience different attenuations.

We present our spatial variation study, where we use both AV and AV500. Figure 3.3 provides the available UDP throughput of single links as a function of the *cable* distance between the source and the destination of the traffic from a single experiment. There is a clear degradation of throughput as the distance increases. However, because of the diversity in positions and types of connected appliances, there is a large range of possible throughputs at any specific distance. We observe that small distances (<30 m) guarantee good links, but that large distances (30–100 m) can yield either good or bad

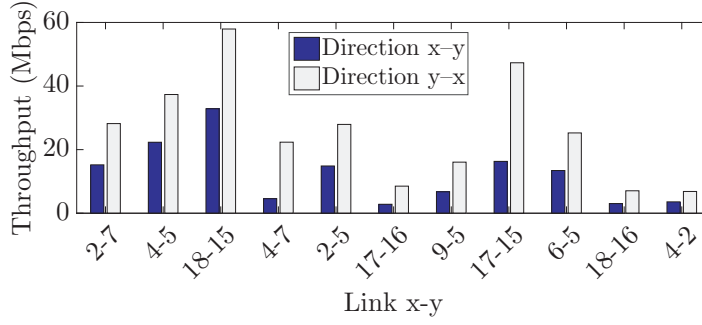


Figure 3.2 – Available-throughput asymmetry in PLC links.




Figure 3.3 – Available throughput vs. cable distance between source and destination for all links of the testbed (left). PB_{err} vs. available throughput for AV (right).

links. By comparing AV and AV500, we observe that AV500 enables some links with no AV connectivity to still enjoy a non-zero throughput, but with severe asymmetries (e.g., link 10-2 with 10x asymmetry).

To further explore the causes of PLC attenuation and variability, we run experiments with two stations connected by a long electrical cable and without any devices attached. We notice that the attenuation in an up to 70-m cable causes a throughput drop of at most 2 Mbps. The attenuation is therefore caused by the multi-path nature of the channel, as explained in Section 2.2. By plugging electrical appliances in this isolated test, we observe that asymmetry was introduced, as also found by Murty et al. [24].

From the above, PLC spatial variation depends on diverse factors: (i) the structure of the electrical network, i.e., the distribution boards; (ii) the electrical appliances attached and their position on the grid; and (iii) (weakly) the cable-distance between stations. The channel is very asymmetric and this is a key feature for spatial variability.

To optimize performance not only in terms of throughput but also delay, hybrid networks need some estimation of the retransmissions a frame suffers due to channel errors. We also evaluate the relationship of the metric PB_{err} with the available throughput.

Figure 3.3 illustrates PB_{err} versus the available throughput for all the links of our testbed. It shows that PB_{err} decreases as throughput increases, as expected. However, because the tone maps are updated based on this metric, some average links might have PB_{err} lower than the best links of the testbed. We further study the PB_{err} metric in Section 3.5, by delving into packet retransmissions. We show that PB_{err} can be used to predict the expected number of retransmissions due to errors.

3.3 Temporal Variation of PLC

We discuss the timescales within which the channel varies. These timescales were introduced for channel modeling and simulation by Sancha et al. [20] (from which we borrow the terminology to name these timescales).

We examine separately the two main components of channel modeling, i.e., the variation of noise generated by the attached electrical appliances and the variation of channel transfer-function (or attenuation). We employ *BLE* to investigate the main properties of PLC channels, by using existing commercial devices. We show that *BLE* reflects the channel quality and the fundamental features of PLC channel modeling explained in Section 2.2. We first focus on noise generated by electrical appliances. It has been shown by measurements [17] that the noise level varies across subintervals of the mains cycle, which yields the first scale governing PLC temporal variation (scale (i)). Due to the periodic nature of the mains, this noise also varies in a scale of multiples of the mains cycle, which results in another timescale for the temporal variation (scale (ii)).

We next focus on attenuation. As discussed in Section 2.2, attenuation is introduced due mainly to impedance mismatches in the transmission line (electrical cable), which are created by connected appliances. As expected, this attenuation changes when the structure of the electrical network changes, hence on a scale of minutes or hours (scale (iii)). This variability strongly depends on appliances usage and on switching the appliances, as this creates aperiodic impulsive noise in the channel.

As hinted above, our study adopts an analysis of three timescales; it is validated by our measurements in the following subsections. Our work differs from [20] in that we examine the channel quality from an end-user and practical perspective, exploring metrics that affect the end-to-end performance. The three timescales are as follows.

- (i) **Invariance Scale:** sub-intervals of the mains cycle, such as the tone-map slots;
- (ii) **Cycle Scale:** multiples of the mains cycles—depends on the noise produced by appliances;
- (iii) **Random Scale:** minutes or hours—related to connection or switching of electrical appliances and depends on human activity.

We introduce our variables, starting with some notation. For the invariance scale, we use the term tone-map slots for the subintervals of the mains cycle, as we can measure the channel quality with respect to tone-map slots by using PLC devices. Let L be the




Figure 3.4 – BLE (capacity) temporal variation.

total number of tone-map slots of the mains cycle, with each slot s having a duration T_s , so that the total slots duration $\sum_{s=1}^L T_s$ is equal to a half mains period (as specified in [10]). Let BLE_s , $1 \leq s \leq L$, denote the BLE of tone-map slot s . In order to study the channel with respect to the three scales defined above, we assume that time t is discrete, with one time unit having real-time duration equal to the mains cycle.

The process $BLE_s(t)$ is different for each link and its distribution can be time-varying over the random scale for a specific link, due to the different types of operating appliances and to different channel transfer-functions. The exact characterization of $BLE_s(t)$ is out of the scope of this work. In our study for cycle-scale variation, we study how often the value of $BLE_s(t)$ changes and how its variance behaves with respect to the link quality. Figure 3.4 illustrates the three timescales and the factors that cause variability. We next examine each timescale.

3.3.1 Invariance Scale

The invariance scale of BLE is affected by the noise levels that appliances produce at different sub-intervals of the mains cycle, and it has direct consequences on estimating link metrics. All our tests show that noise has varying levels over different tone-map slots, thus confirming in practice the existence of periodic impulsive noise, synchronous to the mains (see Section 2.2). Figure 3.5 shows the instantaneous BLE_s from captured frames in typical examples of good and average links. We observe that in HomePlug AV, the total duration of the six tone-map slots is equal to half of the mains cycle, thus BLE_s change periodically, with a period of 10 ms. Each PLC MPDU uses a different BLE_s , depending on which tone-map slot s its transmission takes place.

We highlight that this timescale is crucial for capacity estimation in PLC. With the




Figure 3.5 – Invariance-scale variation of BLE from captured PLC frames of two representative links.

examples of Figure 3.5, we observe that there might be significant variations along the mains cycle, even for good and average links. Thus, link metrics have to be estimated or averaged over all $L = 6$ tone-map slots of the periodic modulation scheme of PLC.

3.3.2 Cycle Scale

We examine the average time during which the quality of the links is preserved in the cycle scale. This sheds light on the average length of probing intervals for link metrics, as there exists a tradeoff in probing: too large intervals might yield a non-accurate estimation, whereas too small intervals can generate high overhead.

We conduct experiments that last 4 minutes, over all links of the testbed. During each experiment, we request BLE_s , $1 \leq s \leq L$, every 50 ms, as this is the fastest rate at which we can currently send MMs to the PLC chip. As we need to avoid random changes in the channel due to switching electrical appliances, all the experiments of this subsection are conducted during nights or weekends. For the cycle-scale variations of the channel, we assume that the electrical network structure is fixed.

Here, we evaluate the cycle-scale variation by using $BLE \doteq \sum_{s=1}^6 BLE_s / 6$, in other words the average BLE over all tone-map slots. We compare the performance between good and bad links¹. Figure 3.6 presents the variation for typical good and bad links of our testbed. Observe that depending on their quality, links exhibit different behaviors. Our findings, validated not only by the representative examples shown here, but also by experiments over an one-year period in all the links of our testbed, are as follows.

Bad Links: Bad links, e.g., 11-4 and 6-5, tend to modify the tone maps much more often than good links do. Moreover, they yield a standard deviation of BLE significantly higher than good links.

¹The classification of the links based on their capacity depends on the PLC technology thus, we do not introduce strict thresholds for this characterization.




Figure 3.6 – Examples of BLE cycle-scale variation for links of diverse qualities.

Average Links: Average links, e.g., 18-15 and 1-2, vary less often than bad links, and might preserve their tone maps for a few seconds. During periods when average links vary often, the standard deviation of BLE can be high, depending on the channel conditions.

Good Links: The tone maps of good links can be valid for several seconds, e.g., link 15-18. Good links, such as link 3-1, that often update the tone maps have insignificant increments or decrements, e.g., of up to 1%, or have impulsive drops of *BLE*, e.g., of up to 5%, with the channel-estimation algorithm needing a few time steps to converge back to the average *BLE* value.

Asymmetry in Temporal Variability: By observing links 15-18 and 18-15, we find that the asymmetry discussed in Section 3.2 translates not only in an average-performance asymmetry, but also in a temporal-variation asymmetry.

Channel-Estimation Algorithms: The temporal variation of link 15-18 is similar with AV and with AV500. By noticing the impulsive *BLE* drops in link 18-15 and by comparing AV with AV500, we detect a feature of the channel-estimation algorithm that might be vendor specific: The AV500 performance oscillation shows that, when bursty errors occur, the estimation algorithm returns very low values of *BLE*. This uncovers that temporal variation in PLC link quality also depends on the channel-estimation algorithm and future work should focus on comparing link-metric estimations for different vendors and technologies.

We next corroborate the above findings over all links of our testbed. Let α be the inter-arrival time of two consecutive *BLE* updates. Figure 3.7 shows the results of the average value of α and the standard deviation of *BLE* for all links sorted by increasing *BLE* order, i.e., link quality. We observe that good links tend to update less often their tone maps, and that *BLE* variability is smaller, compared to bad links. Although some good links might update *BLE* at a similar frequency as bad links (~ 100 ms), as we discussed above, these links tend to have small increments and decrements of *BLE*, thus yielding a stable average performance over minutes and a low *BLE* standard deviation.




Figure 3.7 – Cycle-scale variation of *BLE* with respect to the link quality (links are sorted with increasing average *BLE* order).

To reduce overhead, in cycle scales, that is seconds or minutes, good links should be probed less often than bad links. The cycle-scale variation unveils how link metrics should be updated depending on their quality. It also shows that PLC links are quite stable, providing at most 6 Mbps standard deviation over short timescales.




Figure 3.8 – Random-scale variation of PLC over a total duration of 2 days. Metrics are averaged over 1 minute intervals. Every night at 9pm, all lights are turned off in our building, leading to a channel change for PLC.

3.3.3 Random Scale

In Section 3.3.2, we observe that during timescales of seconds, PLC does not vary much, with a standard deviation of throughput up to 6 Mbps. We now look at longer timescales, i.e., in terms of minutes and hours, with two goals: (i) to examine whether some links could be probed at a slow rate, thus reducing overhead, and (ii) to characterize the variability of PLC performance in presence of high and low electrical-loads. To study the channel quality variation over the random scale, we run tests over long periods, i.e., two days and two weeks, for various links. During these tests, we measure throughput, BLE , and PB_{error} every second. We now denote by μ the mean of $BLE = \sum_{s=1}^6 BLE_s / 6$, and by σ its standard deviation.

Figures 3.8–3.10 show the results of our measurements. Our observations are as follows.

Link Quality vs. Time: The variation of μ is governed by the electrical load. The larger the number of switched-on devices is (e.g., at working hours) the larger the attenuation is, and the lower μ is, as discussed in Section 2.2.

3.3. Temporal Variation of PLC

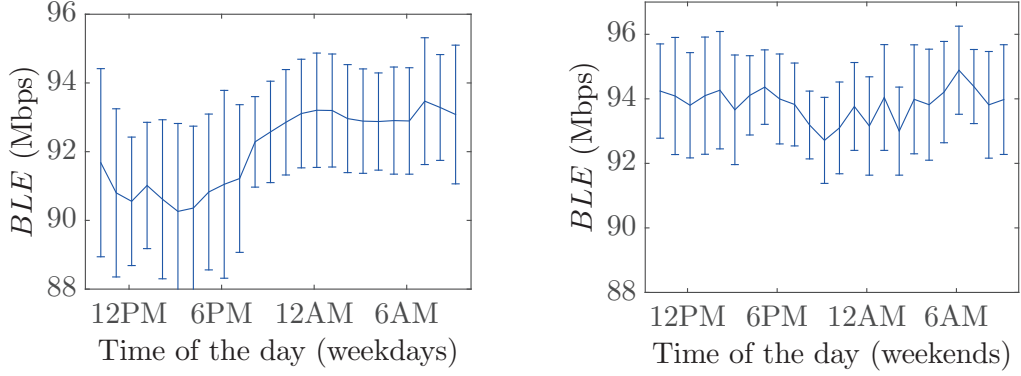


Figure 3.9 – Random-scale variation of BLE for link 1-8 over 2 consecutive weeks. Lines represent the BLE averaged over the same hour of the day and error bars show standard deviation. Station 8 is located in our lounge, hence link quality is affected by the operation of household appliances, such as microwave oven and fridge.

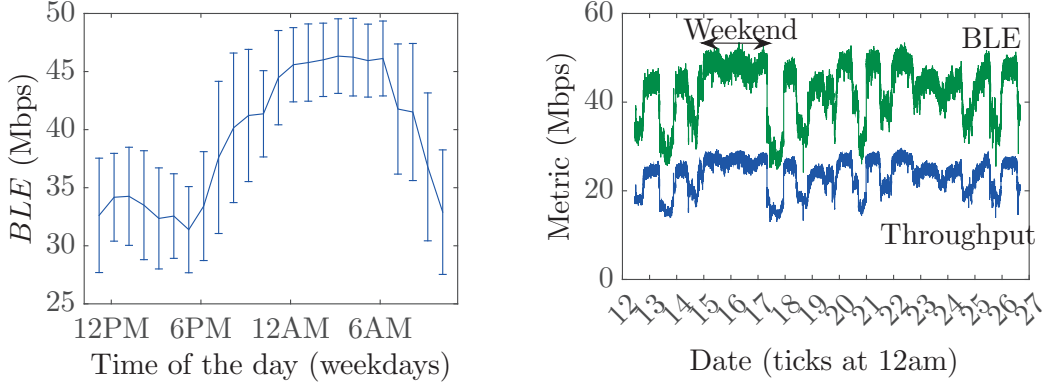


Figure 3.10 – Random-scale variation of BLE for link 2-11 over 2 consecutive weeks in November 2014. The link quality is stable during the first weekend, but not during the second, due to an event held in our building and to high electrical activity.

Link Quality vs. Variability: In Figures 3.9 and 3.10 that represent a good and a bad link, observe the differences in the y-axis scales. For a given link, the random-scale variation of σ strongly depends on the noise of the electrical devices attached, and it is higher when μ is lower. High σ and low μ imply that more devices are switched on therefore, more noise is produced. They also imply that devices are switched on/off more often, creating impulsive noise phenomena. σ is very small for good links; it increases as the link quality, i.e., μ , decreases.

Link Probing: Good links exhibit a negligible standard deviation, which implies that they can be probed every minute or hour, depending on the time of the day.

To conclude, PLC links have low temporal-variability. Good and average links are quite stable over time hence, they can reliably serve bandwidth-hungry applications. The quality of bad links has high variance in the random scale. However, the random-scale variation of PLC can be predicted by learning the electrical-load patterns.

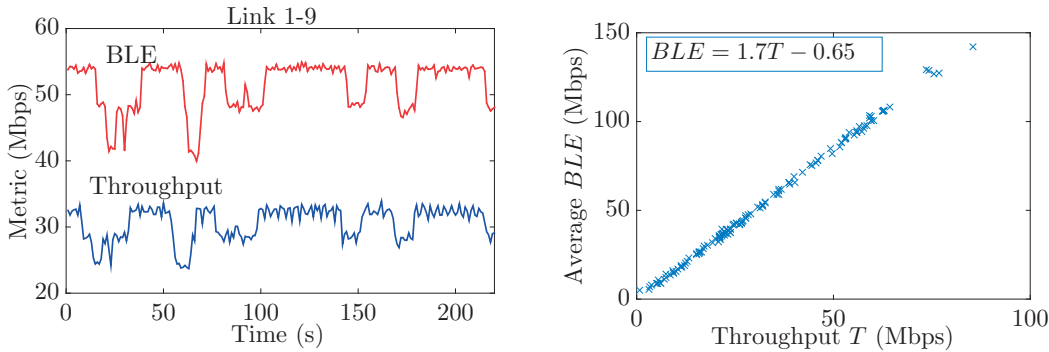


Figure 3.11 – BLE and throughput averaged every 1 s, vs. time, for link 1-9 (left). Average BLE vs. throughput for all the links (right).

3.4 Capacity-Estimation Process

We explore a capacity-estimation process for PLC. As mentioned in Section 2.4, stations estimate a tone map, hence also BLE , if and only if they have data to send. Thus, to estimate link metrics, a few unicast probe packets have to be sent. In the previous section, we discuss how fast the capacity changes given the link quality by sending saturated traffic. Here, we examine how capacity can be estimated with a few probe packets and we explore the size and other features of these packets.

3.4.1 The BLE as a Capacity Estimator

First, we show that BLE , which is included in the frame control of every PLC MPDU, accurately estimates the capacity of any PLC link. We repeat saturated tests for our 144 links and with a duration of 4 min. Figure 3.11 presents the measured throughput and BLE for a sample link (1-9) as a function of time. It also presents average BLE vs. throughput for all links. For each link, we average BLE over all MPDUs transmitted per link. We observe that BLE is an exact estimation of the actual throughput received by the application. Let T be the average throughput. Fitting a line to the data points, we get $BLE = 1.7T - 0.65$. We verified that the residuals are normally distributed.

We next discuss a capacity-estimation technique that uses BLE and probe packets. To conduct a capacity estimation using BLE , a few packets per mains cycle and per estimation interval should be captured, given our temporal variation study in Section 3.3. Here, we investigate an alternative technique that uses MMs to request the instantaneous BLE . The PLC devices provide statistics of the average BLE used over all six tone-map slots. Probe packets do not need to be sent at all tone-map slots of the mains cycle because, according to IEEE 1901 standard, the channel estimation process yields a BLE for all slots when at least one data-packet is sent (see Section 2.4.2).

We now explore whether the number of the probes affects the estimation. To this end, we reset the devices before every run. We perform experiments to estimate the

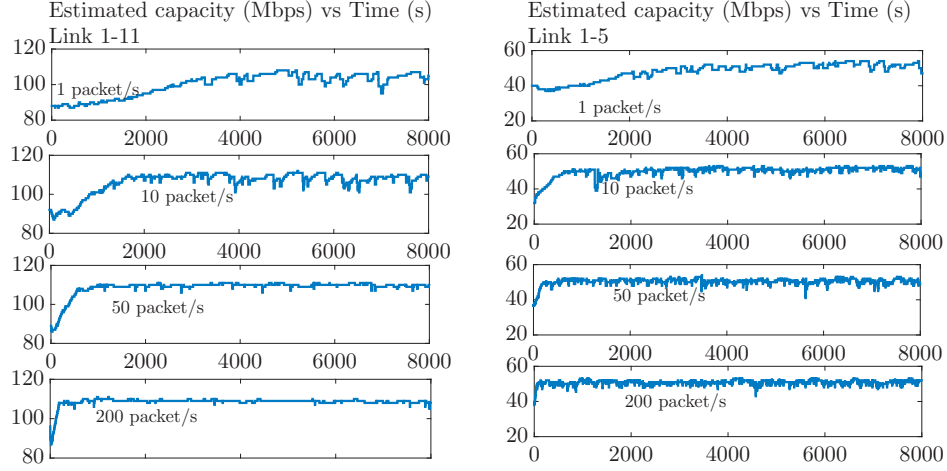


Figure 3.12 – Estimated capacity with different number of packet-probes per second.

capacity, by sending only a limited number of packets of size 1300 B per second (1–200 packets per second)². Figure 3.12 shows that the estimated capacity converges to a value that does not depend on the number of packets sent; however, the number of packets sent per second affects the convergence time of the real estimation. We observe that the channel-estimation algorithm can have a large convergence-time to the optimal allocation of bits per symbol for all the carriers, because it needs many samples from many PBs to estimate the error for every frequency, i.e., carrier. This convergence time depends on the (vendor-specific) channel-estimation algorithm and on the initial estimation (which was reset by us).

To evaluate the convergence time also in realistic scenarios, we perform a test in which we reset the devices at the beginning; but after 2 minutes we pause the probing for approximately 10 minutes. We observe that the PLC devices maintain the channel-estimation statistics, as the estimated capacity resumes from the previous value before stopping the probing process. In tests with AV500 devices, we observe that the capacity estimation starts from a very high *BLE* value, hence it converges faster to lower values for any link quality. Therefore, the convergence time of the capacity estimation by using *BLE* does not necessarily apply in realistic probing conditions.

To conclude, capacity should be estimated by sending probe packets and measuring *BLE* in PLC networks. To estimate capacity, given our study in Section 3.3.1, we take into account the invariance scale and either compute the average $BLE = \sum_{s=1}^6 BLE_s/6$, by capturing the PLC MPDUs or request it using MMs.

3.4.2 Size of Probe Packets

We now investigate the size of probe packets. We observe that for the special case of sending one probe-packet with size less than one PB per second, the estimation might

²Probe packets can be of any size. PLC always transmits at least a PB (512 B), using padding.




Figure 3.13 – Estimated capacity for link 11-6: 1 probe packet per sec., various sizes.

converge to a value smaller than the true one for AV and remain constant with time, independently of the channel conditions.

A representative example of this behavior is shown in Figure 3.13, where AV capacity converges to approximately 89 Mbps when sending only one packet per second with size less than one physical block (520 B including PB header of 8 B). After this convergence, the estimated capacity remains constant. A simple computation shows that the rate required to transmit one PB in one OFDM symbol is $R_{1sym} = (520 \times 8)/T_{sym} \approx 89.4$ Mbps with AV, given the symbol duration $T_{sym} = 40.96 \mu s$. When sending packets smaller than one PB, the rate converges to R_{1sym} for all six slots of the mains cycle, because increasing the rate does not reduce the transmission time (it is not possible to transmit less than one OFDM symbol) while decreasing the probability of error (higher rates yield less robust modulation schemes). Hence, we find that to estimate the capacity of a link by sending only one probe packet per second, it is crucial to send packets larger than one PB or one OFDM symbol.

3.4.3 Priority Class of Probe and Control Packets

While running our experiments with different priority classes (see Section 2.5.2), we noticed an abnormal behavior with priority class CA3 and UDP traffic: There are periodic throughput drops that are not due to *BLE* for all the links. An example of such behavior is depicted in Figure 3.14. *BLE* is an accurate estimation of throughput for the same link and class CA1 in Figure 3.11. However, for CA3 class we observe throughput drops periodically. By using `tcpdump`, we discover that the throughput drops are due to the fact that address resolution protocol (ARP) packets have a priority (CA1) lower than the data packets (CA3), thus they do not contend for the medium, unless the CA3 priority queue is empty. As explained in Section 2.5.2, stations must defer their transmission in case of a higher priority class in the network based on the priority-resolution process. While ARP packets await for the CA3 queue to be emptied, new data-packets are discarded, as the ARP cache does not contain any information about the IP address of the destination.

We observe that all links exhibit the same behavior with the CA3 class. Figure 3.14 also presents a histogram of the inter-arrival times of throughput drops of all links of




Figure 3.14 – (Left) *BLE* and throughput of CA3 class averaged over 1 s intervals, vs. time. (Right) Histogram of inter-arrival times of CA3 throughput drops from 90 links.

our testbed. Let b be the `base_reachable_time` the Linux system uses [25]. Then, the ARP requests are sent when the entry for the neighbor expires in the ARP cache, which happens periodically with a period chosen uniformly at random over the range $[b/2, 3b/2]$. b defaults to 30s for our boards. We observe that the values of the histogram in Figure 3.14 lie in this interval. We do not observe this phenomenon with TCP traffic, because the TCP-ACKs sent by the destination station are used to update the ARP table by our boards. This is obviously a crucial feature of PLC implementations, which can harm the quality-of-service of applications that are tagged with the CA2 or CA3 classes.

Given these experimental results, in hybrid network implementations, management or control messages should be tagged as high priority. Control messages or probe packets, such as ARP, should be tagged with the same priority as the data frames.

3.4.4 Validation of our Capacity Estimation in Hybrid Single-Hop

To further validate our capacity-estimation method, we employ a simple load-balancing algorithm that aggregates bandwidth between WiFi and PLC and operates between the IP and MAC layers. To implement our algorithm, we use the *Click Modular Router* [26]. We forward each IP packet to one of the mediums with a probability proportional to the capacity of the medium. At the destination, we re-order the packets according to a simple algorithm that checks the identification sequence of the IP header.

To estimate the capacities, we probe links with one packet per second. The capacity for PLC is estimated using *BLE*, i.e., averaged over the six tone-map slots of the invariance scale; whereas for WiFi, the capacity is estimated using the modulation and coding scheme index (MCS)³. MCS is averaged over the transmissions (data and probes) every second because, as we observe in Section 2.7.1, WiFi varies more than PLC within a second. Our load-balancing algorithm takes into account our temporal-variation study

³Given MCS index, we retrieve the capacity from IEEE 802.11n standard tables [27].




Figure 3.15 – Performance boost by using hybrid WiFi/PLC, and our load-balancing and capacity-estimation techniques.

on PLC: In Section 3.3.1, we uncover that the PLC channel quality is periodic, with every packet using a different *BLE*. Because an accurate synchronization at this time-scale is challenging for algorithms operating above the MAC layer (such as in IEEE 1905 standard), the capacity of PLC in hybrid networks has to be estimated by averaging over the invariance scale.

In Figure 3.15, we first present the throughput of experiments on one link. We run four experiments back-to-back: using only one of the interfaces (WiFi, PLC) in two; using both interfaces and our load-balancing algorithm (Hybrid) in one; and using both interfaces and a round-robin scheduler for the packets (Round-robin) in the last one. We observe that by using simple load-balancing and reordering algorithms, and our capacity-estimation technique, we can achieve a throughput that is very close to the sum of the capacities of both mediums. In contrast, the throughput of a round-robin scheduler, which has no information on capacity, is limited to twice the minimum capacity of the two mediums (i.e., WiFi in this example), because it assigns the same number of packets to each medium and the slowest medium becomes a bottleneck. To evaluate our algorithm across our testbed, we also compare the completion times of a 600-Mbyte file download by using (i) only WiFi, and (ii) Hybrid⁴. We observe in the same figure, a drastic decrease in completion times when using both mediums.

Our tests validate our capacity estimation methods. They also show that, to exploit each medium to the fullest extent, accurate link-quality metrics are required. However, an open question to be answered is: How should the link metrics be updated to take into account delay or contention? In the next section, we investigate another link metric, i.e., the expected number of retransmissions and the performance of link metrics with respect to background traffic.

⁴Contrary to WiFi, PLC uses queues that are non-blocking: the transport layer is not stopped from sending packets when the MAC queues are full. For these experiments, we omit PLC tests, as dropped packets yield an unfair comparison. In the hybrid test, we send packets at a rate slightly lower than the aggregated capacity to avoid packet losses.

3.5 Retransmitting in PLC Channels

Capacity is a good metric for link quality. However, it does not take into account interference, which is very important for selecting links with high available bandwidth. Moreover, another metric could be useful for delay-sensitive applications that do not saturate the medium but have low delay requirements. Delay is affected by retransmissions either due to bursty errors or to contention; and metrics, such as PB_{err} introduced in Section 2.4.2 (or *packet errors* of IEEE 1905 [4]), are related to retransmissions. We explore the mechanism of retransmissions in PLC networks. We first study the expected transmission count (ETX). Numerous works, e.g., [13, 14], study this metric (or its variations) in WiFi networks by sending broadcast probes. We examine how ETX performs in PLC and the relationship between broadcast and unicast probing.

After studying retransmissions due to errors, we evaluate the sensitivity of link metrics to background traffic. Link metrics in hybrid networks should estimate the amount of background traffic, or be insensitive to background traffic [14]. Thus, a critical challenge is to design link metrics that achieve one of the aforementioned properties.

3.5.1 Retransmission Due to Errors

We first explore how ETX would perform in PLC by sending broadcast packets. Because broadcast packets are transmitted with the most robust modulation and are acknowledged by some proxy station [10], we expect that this method yields very low loss rates.

For the purpose of this study, we set each station in turn to broadcast 1500-byte probe-packets (one every 100 ms) for 500 sec. The rest of the stations count the missed packets by using an identification in our packet header. We repeat the test for all stations of the testbed during night and working hours (day). Figure 3.16 shows the loss rate as a function of throughput and PB_{err} , from our tests for all station pairs. Each pair is represented with its link throughput (respectively, PB_{err}) during the night experiment. The loss rate of broadcast packets in PLC is a very noisy metric for the following reasons:

(i) A wide range of links with diverse qualities have very low loss rates ($\sim 10^{-4}$), and some links even have 0 loss rates. By observing high loss rates, e.g., larger than 10^{-1} , ETX can classify bad links in PLC; but nothing can be conjectured for link quality from low loss rates.

(ii) There is no obvious difference between experiments during the day, when the channel is worse, and night. A few bad links have poor loss rates during the day, but at the same time, a few average links yield even lower loss rates.

(iii) As PLC adapts the modulation scheme to channel conditions when data is transmitted, broadcast packets—sent at the most robust modulation scheme—cannot reflect the real link quality. Given the low loss rates of a wide range of links, ETX appears to be 0 at short-time scales, which provides no or misleading information on link quality.




Figure 3.16 – Loss rate for broadcast packets vs link throughput and PB_{err} for all station pairs.

Due to the above observations, we further explore the mechanism of retransmissions with respect to link quality with unicast traffic. We now delve into the retransmissions of PBs by sending unicast, low data-rate traffic, i.e., 150 Kbps, and by capturing the PLC frame headers. Under this scenario, an Ethernet packet of 1500 bytes is sent approximately every 75ms. The test has a duration of 5 minutes per link. As we discuss above, broadcast packets might be missed by some stations when channel conditions are bad, because they are not retransmitted as soon as a proxy station acknowledges them. In contrast, unicast packets are being retransmitted until the receiver acknowledges them, hence they are always received. For this reason, we look at the frame header (SoF) to study retransmissions. Because there is no indication on whether the frame is retransmitted in the PLC SoF, we employ the arrival time-stamp of the frame to characterize it as a retransmission or new transmission (if the frame arrives within an interval of less than 10ms compared to the previous frame, then it is a retransmission). We also measure PB_{err} every 500 ms.




Figure 3.17 – U-ETX vs. BLE and U-ETX vs. PB_{err} .

We conduct the experiment described above for all the links of our testbed. We compute the unicast ETX (U-ETX) for all the links of the testbed. We count the total number of retransmissions for a packet of 1500 bytes, which produces three PBs. A retransmission occurs if at least one of these PBs is received with errors. Figure 3.17 presents U-ETX as a function of average BLE (with links sorted in increasing BLE order)

and PB_{err} . U-ETX is measured by averaging the number of PLC retransmissions for all packets transmitted during the experiment. We also plot error-bars with the standard deviation of the transmission count. We observe that link quality is negatively correlated with link variability, a conclusion made also when exploring BLE in Section 3.3.2. The higher the U-ETX is, the higher the standard deviation of transmission count is. Links with high BLE are very likely to guarantee low delays, as U-ETX does not vary a lot. U-ETX and the averaged PB_{err} are highly correlated, with almost a linear relationship.

3.5.2 Retransmission Due to Contention

To explore the sensitivity of link metrics to background traffic and to examine how interference can be considered in link metrics, we now experiment with two contending flows. We set a link to send unicast traffic at 150 Kbps, as in the previous subsection, thus emulating probe packets. After 200 seconds, we activate a second link sending “background” traffic at various rates. We measure both BLE and PB_{err} . In these experiments, we observe that BLE is insensitive to low data-rate background traffic for all pairs of links. However, BLE appears to be affected by high data-rate background traffic on a few pair of links. So far, we have not found any correlation between these pairs of links. We explain this phenomenon with the “capture effect”, where the best link decodes a few PBs even during a collision due to very good channel conditions, thus yielding high PB_{err} . In this case, the channel-estimation algorithm cannot distinguish between errors due to PHY layer and errors due to collisions, hence it decreases BLE . Figure 3.18 presents two representative examples of link pairs for which BLE is sensitive and nonsensitive to high data-rate background traffic. Observe that PB_{err} explodes in link 6-11, which is sensitive to background traffic.

To address the sensitivity of our link-metrics to high data-rate background traffic, we take advantage of the frame aggregation procedure of the MAC layer, which is described in Section 2.5. We observe that transmitting a few PBs per 75 ms (150-Kbps rate) yields a sensitivity of metrics to background traffic. However, when two saturated flows are activated, we never notice an effect on BLE . Due to frame aggregation, packets from different saturated flows have approximately the same frame length (i.e., maximum) and when they collide, the channel estimation algorithm works more efficiently than when short probe-packets collide with long ones. To emulate the long frame lengths of saturated traffic, we send bursts of 20 packets such that the traffic rate per second (i.e., the overhead) is kept the same (150 Kbps). In Figure 3.19, we show another link for which BLE is sensitive to background traffic, and the results of our solution. By sending bursts of probe packets, BLE is no more affected by background traffic. This shows that by exploiting the frame aggregation process⁵, we can solve the problem of sensitivity of link metrics to background traffic.

⁵Depending on the PLC technology, these bursts can be transmitted such that only one PLC frame is generated, hence without large MAC overhead.

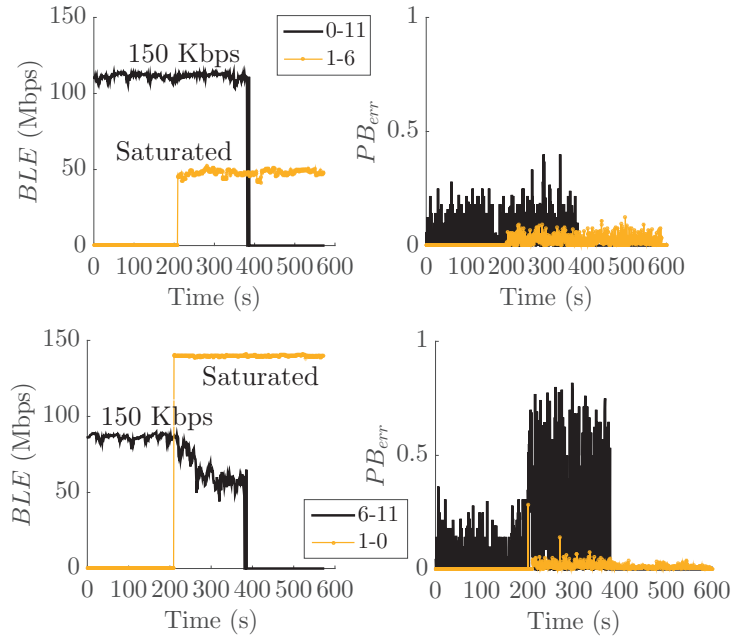


Figure 3.18 – Link metrics of 2 sets of contending links with low data-rate and saturated traffic. The link-metrics on the bottom plot are sensitive to background traffic.

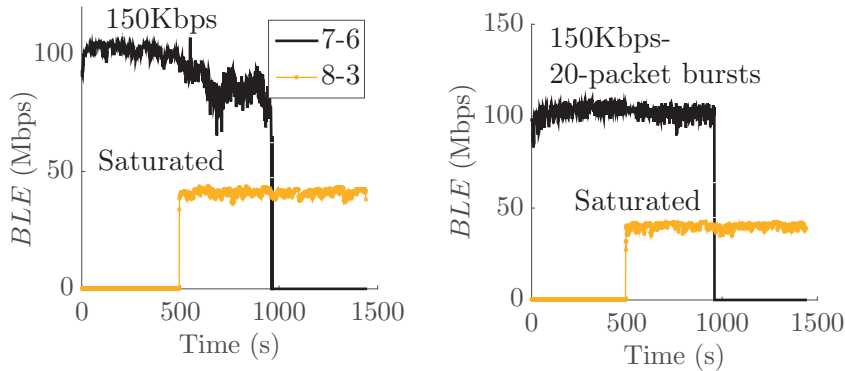


Figure 3.19 – Addressing the link-metric sensitivity to background traffic by sending bursts of probes.

To conclude, we have studied the mechanism of retransmissions in PLC. Although broadcast probe-packets yield significantly less overhead in link-quality estimation, they do not provide accurate estimations. In contrast, unicast probe-packets reflect the real link quality, but by producing more overhead. We observe that PB_{err} can be used to estimate U-ETX and to indicate interference in PLC. However, estimating the amount of interference is challenging and should be further investigated. We leave this extension for future work. We introduce techniques to address potential sensitivity of link-metric estimation to background traffic.

3.6 Link-Metric Guidelines

One of our significant contributions in this chapter is the systematic guidelines for link-metric estimation in PLC networks. In this section, we summarize these guidelines and facilitate the future implementation of heterogeneous networks that include PLC. Table 3.1 outlines our guidelines for efficient link-metric estimation with PLC, given our experimental study in this chapter. We present our link metrics and instructions on how to perform their estimation with high accuracy and low overhead. In addition, we refer the reader to the relevant sections in this chapter.

Policy	Guideline/Explanation	Section
Metrics	BLE and PB_{err} , defined by IEEE 1901.	3.4, 3.5.1
Unicast probing only	Broadcast probing cannot be used, as it does not give sufficient information on link quality.	3.5.1
Shortest time-scale	BLE should be averaged over the mains cycle.	3.3.1
Size of probes	Larger than one PB (or one OFDM symbol) to avoid inaccurate convergence of rate adaptation algorithm.	3.4.2
Frequency of probes	Should be adapted to link quality for lower overhead.	3.3.2, 3.3.3
Burstiness of probes	Can tackle a potential inaccurate convergence of the channel estimation algorithm or the sensitivity of link metrics to background traffic.	3.4.2, 3.5.2
Asymmetry in probing	There is both spatial- and temporal-variation asymmetry in PLC links. This could affect bidirectional traffic, such as TCP, that requires routing in both directions.	3.2, 3.3.2
Priority of control traffic	Control messages should be tagged at the same priority as the data to avoid performance degradation.	3.4.3

Table 3.1 – Guidelines for PLC link-metric estimation.

We would like to note that using our two metrics, our temporal variation study, and our frame-aggregation model of Section 2.5.1, PLC performance can be accurately characterized and simulated, thus reducing the complexity of exact representation of the channel model and the PHY layer.

3.7 Related Work

A large body of work (e.g., [17, 20]) focuses on channel modeling and noise analysis, and very little work, such as [24], investigates PLC performance from an end-user perspective that is the focus of this chapter. Murty et al. [24] explore the performance of HomePlug AV when household devices operate in the network. They observe that switching the

appliances affects significantly the performance and introduces asymmetry, and that different appliances create diverse noise levels. Compared to our work, they do not study the temporal variation of performance and they treat PLC as a black-box, i.e., without any PHY or MAC layer background, due to lack of standardization.

There are many works on PLC-channel characterization that either rely on simulation or channel measurements, but they do not study the performance with respect to higher network layers. For instance, Weling [28] presents measurements of 40.000 PLC channels in order to explore an adaptive transmit-power management. The field tests reveal that the SNR varies during the day, and is higher at night hours where most probably less electrical appliances are used. The author unveils how average throughput or SNR achieved varies across countries.

There are experimental PLC studies that explore both PHY and MAC characteristics. In particular, Liu et al. [29] employ a testbed to investigate the interoperability and coexistence of different HomePlug AV networks and propose a scheme that can be employed to ameliorate the performance of multiple contending AV networks. In addition, they verify that HomePlug AV stations can communicate between different phases and that there is no rate-anomaly phenomenon in PLC, which occurs when a station with a low rate bounds the throughput of all stations and which exists in WiFi networks [30]. Zarikoff and Malone [31] study PHY and MAC features under ideal channel conditions.

A few works focus on comparing the WiFi and PLC performance, e.g., [32, 33]. Lin et al. [32] investigate older specifications of PLC and WiFi, i.e., HomePlug 1.0 and 802.11 a/b, respectively. The authors provide testbed measurements from 20 houses for metrics such as coverage, throughput, and connectivity. Tinnakornsriruphap et al. [33] introduce a comparison between hybrid PLC/WiFi networks and single-technology networks. The authors find that hybrid networks contribute to increased coverage in home networks; they also argue that using alternating technologies for multi-hop routes yields good performance. However, they do not study link metrics that can be used to optimize routing in such networks. Due to the complexity of representing PLC in simulations, the authors plug experimental measurements into their simulator for link capacities.

3.8 Summary

In this chapter, we explore experimentally the end-to-end performance of isolated PLC links in an enterprise environment, focusing on spatial and temporal variations of capacity. We uncover that PLC provides very good connectivity for stations connected to the same distribution panel, forming an almost fully-connected graph. However, the exact link quality is challenging to predict given station location, due to the severe asymmetry PLC links exhibit. Concerning the temporal variation, we find that it occurs in three time-scales because of the diverse types of noise in PLC channels and because of the human behavior with respect to electrical activity. Yet, PLC links have low variability, hence they can serve bandwidth-hungry applications.

We also study PLC link metrics and their variation with respect to space, time, and background traffic. Similar metrics have long been pursued by the research community for WiFi and have been required by the recent standardization of hybrid networks. We introduce practical guidelines on efficient metric-estimation in hybrid implementations. We observe that there is a high correlation between link quality and its variability, which has a direct impact on probing overhead and accurate estimations.

4 Decoupled Analysis and Enhancement of Multi-user Performance

4.1 Introduction

In the previous chapter, we focused on the efficiency and variability of *single* PLC links. In this chapter, we explore PLC performance under *multi-user* settings, where multiple stations contend for the medium. We raise two fundamental questions: (i) How does performance scale with the number of contending stations? (ii) How fast and efficiently does the protocol adapt to dynamic traffic? To answer these questions, we employ an analytical framework for modeling and enhancing multi-user performance, WiFi and PLC testbeds to validate our model and enhancements and finally, extensive simulations.

As we observed in Section 2.5, PLC is a shared medium and it employs a CSMA/CA protocol in order to resolve contention issues among stations. The PLC CSMA/CA includes an additional variable compared to the one of WiFi; it is called the *deferral counter* and introduces high complexity. Here, we show that, despite its intricacies, *the deferral counter brings in high efficiency and adaptability under high-intensity or fast-varying traffic*. In the following, we provide the reasons for these benefits of the PLC CSMA/CA, by comparing it with the one of WiFi. After a review of the two CSMA/CA mechanisms, we describe how the remainder of the chapter is organized.

4.1.1 Background and Motivation: IEEE 802.11 vs. 1901 CSMA/CA

Let us recall the main differences between WiFi and PLC CSMA/CA protocols, hereinafter referred to as 802.11 and 1901, respectively. Figure 4.1 shows the backoff process of both 802.11 and 1901 as a finite-state machine. The differences of 1901 with 802.11 are represented with dashed lines. Observe the main difference between 802.11 and 1901: After sensing the medium idle following a busy period, 1901 stations might double CW and select a new backoff counter BC depending on the deferral counter value DC (see Algorithm (ii)). In contrast, 802.11 stations simply resume their BC . The key parameters that enable 1901 to adapt to traffic demands are the *initial values of the deferral counter* at each backoff stage i , d_i . Remember that a station moves to backoff stage $i + 1$, if it senses $d_i + 1$ transmissions before its BC has expired, by the definition of the protocol.

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance

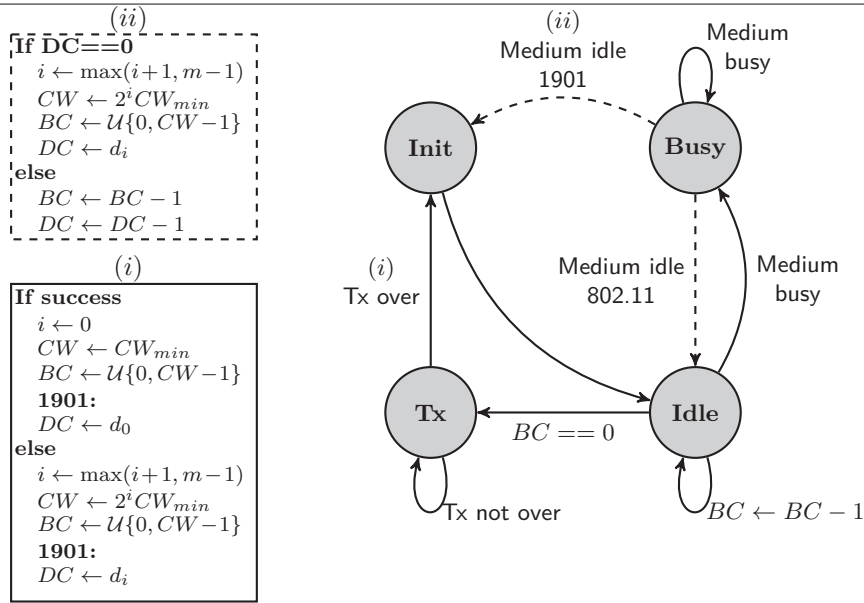


Figure 4.1 – A finite-state machine that represents the 802.11 and 1901 backoff procedures, as running in a station at each time-slot. A station is in state **Init** whenever it selects BC (or updates DC in 1901), in states **Busy** and **Idle** when the medium is sensed busy and idle respectively, and in state **Tx** when it is transmitting. After a transmission attempt, stations move to state **Init** and execute Algorithm (i). In 1901, after choosing a new BC , the station updates DC according to the backoff stage i . The parameter m denotes the total number of backoff stages. The dashed lines highlight the fundamental difference between the two standards. After sensing the medium idle while being in **Busy** state, 1901 stations move to state **Init** and execute Algorithm (ii). In contrast, 802.11 stations return to the **Idle** state and just resume BC .

Let us underline a major consequence of the difference between 1901 and 802.11, which is due to the introduction of the deferral counter in 1901. A 1901 station can redraw its backoff counter *without* experiencing a collision, if its deferral counter is 0 and it senses the medium busy. This results in 1901 having two distinct features that offer new opportunities for access control:

- With 1901, the contention window can be increased without suffering from a collision, hence a station can adapt its transmission behavior to the level of contention in the network, without wasting channel time in collisions.
- By appropriately configuring its deferral counter, 1901 can optimally adjust its level of reaction to contention with a fine granularity, which is contrary to 802.11, where the level of reaction cannot be tuned.

Figure 4.2 provides evidence of the 1901 potential. We investigate the time evolution of the protocol dynamics and compare 1901 with 802.11 for two different values for the number of stations in the network, $N = 2$ and $N = 15$. Observe that for $N = 2$, 802.11 has a large average CW only after consecutive collisions, whereas 1901 adapts its CW




Figure 4.2 – Time evolution of the contention window CW averaged over all stations in the network (left axis) and the binary outcome of transmission attempts (right axis), i.e. success (S) or collision (C), in simulation. All stations are saturated and start at backoff stage 0. 1901 is simulated with CA1 class, and 802.11 with $CW_0 = 16$ and 7 backoff stages (i.e., 802.11a/g/n). Our simulator is validated experimentally in Section 4.4.1, and we use the timings specified in Section 4.4.

even after successful transmission attempts. The 1901 average CW increases or remains constant as a given station keeps transmitting, and decreases only when the other station successfully wins the medium. Finally, notice that for $N = 15$, 1901 adapts its CW before a collision, thus it wastes less time in collisions and the average CW converges much faster to the steady state. Hence, under dynamically changing traffic, 1901 can adjust its CW to the load demand faster than 802.11 can.

In this chapter, we show that the two above-mentioned 1901 features make it possible to configure 1901 so that it efficiently adjusts to the level of congestion in the network, *without knowing the number of active stations*. In contrast, although many approaches for efficiently configuring and optimizing 802.11 have been proposed in the literature, they are limited to a given (known) number of active stations, and 802.11 optimal configurations need to estimate the number of contending stations.

Motivated by the above, we propose a framework for modeling and enhancing the 1901 CSMA/CA process. In Section 4.2, we introduce a model that accurately captures 1901 performance and requires very small state-space; this model comes in the form of a fixed-point equation that we show admits a unique solution for a wide range of configurations. To model the complex 1901 protocol, we rely on the *decoupling assumption* which asserts that the backoff processes of the stations are independent. We discuss the accuracy and the validity of this assumption later in this dissertation. We employ our model to devise a configuration that boosts throughput in Section 4.3. Our configuration consists in simply setting existing MAC parameters to appropriate values, thus it can be readily implemented by PLC manufacturers.

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance

In Section 4.4, we conduct a thorough evaluation of our model and enhancements. The proposed configuration provides significant performance improvements. In the same section, we validate our enhancement in a real testbed, by implementing the 1901 mechanism on WiFi hardware. We review related work on 802.11 and 1901 MAC layers in Section 4.5 and we summarize the results of this chapter in Section 4.6.

One of the most remarkable results of this chapter is that, with the proposed configuration, 1901 provides a performance very close to that of an optimally configured MAC protocol without knowing the total number of contending stations N . In contrast, similar methods for enhancing the 802.11 CSMA/CA process do require knowing N (see, e.g., [34]), which challenges their practicality in real deployments where N varies in time. Thus, with our proposed configuration, 1901 represents an interesting step towards a practical CSMA/CA protocol that performs close to optimum.

4.2 Analysis

We first present the assumptions under which we study the network. These assumptions are adopted in the rest of this dissertation, unless otherwise stated. Next, we discuss the modeling assumption of this chapter. After outlining our assumptions, we introduce our model for the 1901 CSMA/CA protocol.

4.2.1 Network Assumptions

We analyze the protocol under the following assumptions, all of them widely used for modeling 802.11 and 1901 [15, 30, 35, 36].

Perfect sensing: There is a single collision-domain with N stations and there are no hidden terminals.

Saturated queues: All stations are saturated (always have a packet to send).

Perfect channel: There is no packet loss or errors due to the physical layer, and transmission failures are only due to collisions. The stations use the same physical data rate, hence all frames have equal durations¹.

Infinite retry limit: The stations have an infinite retry limit; that is, they never discard a packet until it is successfully transmitted².

Homogeneous network: The 1901 standard introduces four different priority classes (see Section 2.5) and specifies that only the stations belonging to the highest contending priority class run the backoff process. In our analysis, we follow this, and we consider a scenario in which all the contending stations use the same set of parameters, i.e., those of the highest priority-class present in the network.

¹As a result, our model and enhancement can be used for different data-rate HomePlug specifications, as they all employ the same CSMA/CA process.

²Contrary to 802.11, 1901 does not specify a retry limit. However, there is a timeout on the frame transmission that is vendor specific. For instance, for the HomePlug AV devices tested in Section 4.4.1, the timeout for CA1 priority frames is 2.5 s, which is very large compared to the maximum frame duration (2.5 ms [10]). Therefore, the infinite retry limit assumption is reasonable.

4.2.2 Decoupling Assumption

We now turn our attention to the modeling assumption that is referred to as the *decoupling assumption* (see e.g., [15, 30]). According to this assumption, *the backoff process of a station is independent of the aggregate attempt process of the other $N - 1$ stations*. This yields the following approximations:

1. Given a tagged station, the probability that at least one of the other stations transmits at any time slot is fixed, denoted by γ , from which (i) transmission attempts experience a fixed collision probability γ , and (ii) a station with $BC \neq 0$ senses the medium busy at any time slot with a fixed probability γ . With this, we can compute the probability that a station transmits in a randomly chosen time slot, which we denote by τ , as a function of γ .
2. The transmission attempts of different stations are independent, with the transmission probability of a station given by the average attempt rate τ . This enables us to express γ as a function of τ , leading to a fixed-point equation.

4.2.3 The Decoupling-Assumption Model

We are now ready to introduce our model. Our analysis requires computing the expected number of time slots spent by a station at backoff stage i (where a time slot can either be idle or contain a transmission). Let k denote the value of BC drawn uniformly at random in $\{0, \dots, CW_i - 1\}$, when the station enters stage i . If the station is running 802.11, the station leaves the backoff stage when (and only when) it attempts a transmission; hence, the station stays in the backoff stage exactly $k + 1$ time slots, until BC expires and it attempts a transmission. In contrast, in 1901 a station might leave backoff stage i either because of a transmission attempt, when BC expires (like in 802.11), or because it has sensed the medium busy $d_i + 1$ times, before BC has expired. In the latter case, the station spends a number of slots at backoff stage i equal to j , if it senses the medium busy for the $(d_i + 1)$ th time in the j th slot (where $d_i + 1 \leq j \leq k$).

Let us write bc_i for the expected number of time slots spent by a station at backoff stage i . To compute bc_i we need to evaluate the probability of the events that (i) a station attempts a transmission or (ii) senses the medium busy $d_i + 1$ times within the k slots (i.e., before BC expires). Let b be the random variable denoting the number of busy slots within the k slots. Because of the decoupling assumption, b follows the binomial distribution $\text{Bin}(k, \gamma)$. Now, let x_k^i be the probability that a station at backoff stage i jumps to the next stage $i + 1$ in k or fewer time slots due to event (ii). Then,

$$x_k^i = \mathbb{P}(b > d_i) = \sum_{j=d_i+1}^k \binom{k}{j} \gamma^j (1 - \gamma)^{k-j}. \quad (4.1)$$

We can compute bc_i as a function of γ via x_k^i . We distinguish two cases on k . First, if $k > d_i$, then event (i) occurs with probability $1 - x_k^i$, and event (ii) occurs with

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance

probability x_k^i . For event (i), the station spends $k + 1$ slots in stage i . For event (ii), the station spends j slots in backoff stage i with probability $x_j^i - x_{j-1}^i$ for $d_i + 1 \leq j \leq k$ (observe that $x_j^i - x_{j-1}^i$ is the probability that (ii) happens exactly at slot j). Second, if $k \leq d_i$, then event (ii) cannot happen. Thus, the backoff counter expires, event (i) always takes place, and the station spends $k + 1$ time slots in stage i . By considering all the possible cases described above, bc_i can be computed as

$$bc_i = \frac{1}{CW_i} \sum_{k=d_i+1}^{CW_i-1} \left[(k+1)(1-x_k^i) + \sum_{j=d_i+1}^k j(x_j^i - x_{j-1}^i) \right] + \frac{(d_i+1)(d_i+2)}{2CW_i}. \quad (4.2)$$

Next, we need to compute the probability that a station at backoff stage i ends this stage by attempting a transmission, which we denote by t_i , and the probability that such a backoff stage ends with a successful transmission, which we denote by s_i . Similarly to (4.2), t_i can be computed as

$$t_i = \sum_{k=d_i+1}^{CW_i-1} \frac{1}{CW_i} (1-x_k^i) + \frac{d_i+1}{CW_i}. \quad (4.3)$$

From the above, s_i can be simply computed as $s_i = (1-\gamma)t_i$.

Building on the above expressions for bc_i , t_i and s_i , we now address the computation of the average attempt rate of a station, τ . Let R be the random variable denoting the number of transmission attempts experienced by a successfully transmitted packet. Similarly, let X be the random variable denoting the total number of slots spent in backoff for a successfully transmitted packet. After every successful transmission, the station enters backoff stage 0 and an identical procedure is followed for the next packet, independently of previous transmissions. Hence, from the renewal-reward theorem (R being the reward and X the renewal lifetimes [30]), the average attempt rate is given by

$$\tau = \frac{\mathbb{E}[R]}{\mathbb{E}[X]}. \quad (4.4)$$

We compute $\mathbb{E}[R]$ and $\mathbb{E}[X]$ with the following two lemmas, with proofs in Appendix A.1.

Lemma 1. *The expected number of slots spent in backoff per successfully transmitted packet is*

$$\mathbb{E}[X] = \sum_{i=0}^{m-2} bc_i \prod_{j=0}^{i-1} (1-s_j) + \prod_{i=0}^{m-2} (1-s_i) \frac{bc_{m-1}}{s_{m-1}}.$$

Lemma 2. *The expected number of transmission attempts per successfully transmitted packet is*

$$\mathbb{E}[R] = \sum_{i=0}^{m-2} t_i \prod_{j=0}^{i-1} (1-s_j) + \prod_{i=0}^{m-2} (1-s_i) \frac{t_{m-1}}{s_{m-1}} = \frac{1}{1-\gamma}.$$

We can now derive a fixed-point equation on γ . From the decoupling assumption, the probability γ that at least one other station transmits can be expressed as a function of τ as follows:

$$\gamma = \Gamma(\tau) = 1 - (1 - \tau)^{N-1}.$$

In turn, $\tau = \mathbb{E}[R]/\mathbb{E}[X]$ can be expressed as a function of γ by using Lemmas 1 and 2. Let us denote this function by $G(\gamma)$. The composition of the functions $\tau = G(\gamma)$ and $\gamma = \Gamma(\tau)$ yields the fixed-point equation for the collision probability

$$\gamma = \Gamma(G(\gamma)). \quad (\text{FPE})$$

By solving the above fixed-point equation, we can determine the γ value at the operating point of the system. Theorem 1 below establishes the uniqueness of the solution of (FPE) under the condition that the level of aggressiveness (i.e., transmission probability) is decreasing with the backoff stage i . The theorem, with proof in Appendix A.1, provides a wide range of configurations satisfying this condition. From Table 2.2, these configurations are compliant with the standard, except for the class CA2/CA3 at backoff stage $i = 1$. This suggests that it might be worth re-evaluating this configuration choice in the standard: indeed, with class CA2/CA3 stations increase their level of aggressiveness upon jumping to backoff stage 1, which contradicts the spirit behind the protocol design for decreasing aggressiveness upon an indication of congestion.

Theorem 1. $\Gamma(G(\gamma)) : [0, 1] \rightarrow [0, 1]$ has a unique fixed-point if the following condition is satisfied for $0 \leq i \leq m - 2$

$$\begin{cases} CW_{i+1} \geq CW_i & \text{if } d_{i+1} = d_i, \\ CW_{i+1} = CW_i & \text{if } d_{i+1} < d_i, \\ CW_{i+1} \geq 2CW_i - d_i - 1 & \text{otherwise.} \end{cases} \quad (4.5)$$

Throughput Evaluation

We now explain how to obtain actual throughput figures from our model. Once we have the value for γ from the fixed-point equation of (FPE), we can obtain τ . We can then compute p_s and p_e , the probability that a slot contains a successful transmission or that it is empty, respectively, from $p_s = N\tau(1 - \tau)^{N-1}$ and $p_e = (1 - \tau)^N$. We can further compute the probability that a slot contains a collision as $p_c = 1 - p_e - p_s$. We now have enough information to compute the normalized throughput S of the network as

$$S = \frac{p_s D}{p_s T_s + p_c T_c + p_e \sigma}, \quad (4.6)$$

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance

where D is the frame duration, T_s is the duration of a successful transmission, T_c is the duration of a collision, and σ is the time slot duration. We give these parameters in Section 4.4.2.

Under the assumption that the stations are saturated, frame durations are the same for all stations independently of their rates, because of the frame aggregation procedure described in Section 2.5.1 and of the maximum frame duration imposed by the standard. Thus, in order to compute actual throughput figures, we just have to replace the frame duration D with the number of bytes per PLC frame in (4.6). Our analysis aims to provide performance evaluation that is independent of the PLC PHY technology and rates and that can be applied to any HomePlug specification described in Section 2.3. To this end, we employ the normalized throughput given by (4.6).

4.3 Enhancements of IEEE 1901 for High Throughput

As highlighted in Section 4.1.1, 1901 has the advantage over 802.11 in that it reacts pro-actively to collisions; however, the current configuration proposed by the standard does not exploit this advantage to obtain high throughput. In this section, we use our 1901 model to devise efficient configurations that perform close to an optimal MAC protocol.

4.3.1 Deriving the Proposed Configuration

Below, we propose a configuration of the 1901 MAC parameters; it drives the system to the operating point that maximizes the achievable throughput of the network. According to [15], this point is achieved when the transmission probability τ takes the following optimal value, which we denote by τ_{opt} ,

$$\tau_{opt} \approx \frac{1}{N} \sqrt{\frac{2\sigma}{T_c}}. \quad (4.7)$$

From this, we have that the collision probability at the optimal operating point, which we denote by γ_{opt} , is given by

$$\gamma_{opt} = 1 - (1 - \tau_{opt})^{N-1} = 1 - \left(1 - \frac{1}{N} \sqrt{\frac{2\sigma}{T_c}}\right)^{N-1} \quad (4.8)$$

which, as N gets large, can be approximated by³

$$\gamma_{opt} \approx 1 - e^{-\sqrt{\frac{2\sigma}{T_c}}}. \quad (4.9)$$

Following the analysis in Section 4.2, γ can be obtained from the fixed-point equa-

³For the 1901 parameters (shown in Section 4.4.2), the approximation is good for $N \geq 3$.

4.3. Enhancements of IEEE 1901 for High Throughput

tion (FPE), which can be rewritten as

$$1 - (1 - \gamma)^{1/(N-1)} = G(\gamma). \quad (4.10)$$

Figure 4.3 illustrates the solutions of the above equation for (i) 1901 under the default configuration (“1901 default”), (ii) 1901 under the configuration proposed here (“1901 proposed”), and (iii) 802.11 under its default configuration (“802.11 default”), for different values of N ($N = 2, 7, 12, \dots, 27$).




Figure 4.3 – The analysis of our enhancement. The solution to our model is given by the intersection of $\tau = 1 - (1 - \gamma)^{1/(N-1)}$ and $\tau = G(\gamma) = \mathbb{E}[R]/\mathbb{E}[X]$. We plot $1 - (1 - \gamma)^{1/(N-1)}$ for various number of stations N (solid lines) and $G(\gamma)$ for (i) the optimal case on which our analysis relies, (ii) the default 1901 (CA1) configuration, (iii) the default 802.11, and (iv) our proposed enhancement given by (4.16) and (4.17) for $m = 6$. The default 1901 and 802.11 configurations yield solutions with large collision probabilities. In contrast, our proposed 1901 yields solutions much closer to γ_{opt} .

To drive the 1901 network towards the optimal operation point derived above, we need the γ value that solves (4.10) to be as close as possible to $\gamma = \gamma_{opt}$ for all values of N . To force this, we need that $G(\gamma)$ resembles as much as possible a step function that takes very large values for $\gamma < \gamma_{opt}$, decreases sharply to a very small value at $\gamma = \gamma_{opt}$, and takes a very small value from this point on. An example of such an optimal step function is illustrated in Figure 4.3 (“ $G(\gamma)$ optimal”). Note that $G(\gamma)$ is given by $\mathbb{E}[R]/\mathbb{E}[X]$, and its shape depends on the setting of the MAC parameters. Below, we derive the parameter configuration that drives $G(\gamma)$ as close to the desired behavior as possible. To achieve this, we proceed as follows. $G(\gamma)$ can be recast as

$$G(\gamma) = \frac{1}{(1 - \gamma) \sum_{i=0}^{\infty} bc_{\min(i, m-1)} \prod_{j=0}^{i-1} [1 - (1 - \gamma)t_{\min(j, m-1)}(\gamma)]}, \quad (4.11)$$

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance

where t_j is given by (4.3) and bc_i by (4.2). Note that $\min(i, m - 1) = i$, if $0 \leq i \leq m - 1$ and $m - 1$ otherwise. From the above expression, it can be seen that $G(\gamma)$ depends on γ through: (i) the $(1 - \gamma)bc_i$ terms, which decrease with γ , and (ii) the $1 - (1 - \gamma)t_j$ terms, which increase with γ . From the proof of Theorem 1 which is given in Appendix A.1, it can be seen that $G(\gamma)$ strongly decreases with γ , which shows that the main dependency of G on γ comes from the t_j 's. Thus, our goal will be achieved if $t_j(\gamma)$ follows a step function that is as close to 1 as possible for $\gamma < \gamma_{opt}$ and as close to 0 as possible for $\gamma > \gamma_{opt}$: with this, $G(\gamma)$ will take the largest possible value for $\gamma < \gamma_{opt}$ and the smallest possible value for $\gamma > \gamma_{opt}$. The value of t_i is computed as

$$t_i = \frac{1}{CW_i} \sum_{j=0}^{CW_i-1} P(d_i, j, \gamma), \quad (4.12)$$

where $P(d_i, j, \gamma)$ is the probability that a station senses no more than d_i transmissions in j slots, given that the collision probability is equal to γ .

From the analysis of Section 4.2, the number of transmissions that a station senses in j slots follows the binomial distribution $\text{Bin}(j, \gamma)$; hence, its probability mass function (pmf) has a peak around the mean value, $j\gamma$, and has a variance of $j\gamma(1 - \gamma)$. The key approximation that we make to derive the proposed configuration is to assume that the variance is sufficiently small such that $P(d_i, j, \gamma)$ —which is a complementary cumulative distribution function—is equal to 0 when d_i is below the peak (i.e., $j\gamma > d_i$) and equal to 1 otherwise (i.e., $j\gamma \leq d_i$). Note that for typical configurations⁴ we have $\gamma_{opt} \ll 1$, which makes the variance small and thus helps to make this approximation accurate.

With the above approximation, we have

$$t_i \approx \frac{1}{CW_i} \sum_{j=0}^{CW_i-1} 1_{j\gamma < d_i+1} = \frac{d_i + 1}{\gamma CW_i}, \quad (4.13)$$

which (as desired) takes a large value for small γ and a small value for large γ ⁵. To force that the transition from large values (around 1) to small ones (around 0) takes place at $\gamma = \gamma_{opt}$, we require that t_i takes an intermediate value, i.e., $t_i = 1/2$, at this point, which yields

$$t_i|_{\gamma=\gamma_{opt}} = \frac{d_i + 1}{\gamma_{opt} CW_i} = \frac{1}{2}, \quad (4.14)$$

from which

$$d_i = \frac{CW_i}{2} \gamma_{opt} - 1. \quad (4.15)$$

The rationale behind $t_i = 1/2$ is as follows: to force the solution of (4.10) to be at γ_{opt} , we need that the right-hand side of (4.10) to sharply decrease from a large to a small

⁴From (4.9), γ_{opt} decreases with T_c/σ hence, the approximation is good for $T_c/\sigma \gg 1$, which is true for 1901 where $T_c/\sigma \approx 81$ (see the 1901 timing parameters in Section 4.4.2).

⁵Intuitively, (4.13) follows from the assumption that the station transmits, if it selects a backoff counter smaller than the average slots required for the deferral counter to expire, i.e., $(d_i + 1)/\gamma$.

4.3. Enhancements of IEEE 1901 for High Throughput

value at this point, so that the two curves cross at $\gamma = \gamma_{opt}$ for all N . Since increasing γ decreases t_i , if we set t_i such that with γ_{opt} we have $t_i = 1/2$, then for a smaller γ , t_i will be larger and the stations will be more likely to transmit (hence, the transmission probability increases). Otherwise, they will be more likely to defer.

With the above setting of d_i , if the current transmission probability in the network is too high, stations jump to the next backoff stage without transmitting. To ensure this reduces the overall transmission probability, we need to make sure that the CW_i of the next backoff stage is larger than the current one. However, to keep a sufficient level of granularity in the transmission probabilities, the increase of CW_i should not be too drastic (as otherwise a station might jump from a CW_i value that is too small to one that is too large). Following this argument, we set $CW_{i+1} = 2CW_i$. This yields

$$d_i = \left\lceil \gamma_{opt} \frac{2^i CW_{min}}{2} - 1 \right\rceil, \quad (4.16)$$

where the ceiling is used to avoid negative values of d_0 .

From the above, we have fixed the values of d_i . The remaining challenge is to configure CW_{min} , and m . To ensure that the curve $1 - (1 - \gamma)^{1/(N-1)}$ crosses $G(\gamma)$ at $\gamma = \gamma_{opt}$, we need τ_{max} to be sufficiently large so that even with small N the curve is crossed at this point; conversely, we also need that τ_{min} is sufficiently small to cover the cases with large N .

To guarantee that τ_{max} is sufficiently large even for $N = 2$, the CW at backoff stage 0, CW_{min} , needs to be as small as the optimal CW for $N = 2$. Accordingly, we set CW_{min} equal to the optimal CW value for $N = 2$, this gives $\gamma_{opt} = \tau = 2/(CW_{min} + 1)$, from which

$$CW_{min} = \left\lfloor \frac{2}{1 - e^{-\sqrt{\frac{2\sigma}{T_c}}}} - 1 \right\rfloor. \quad (4.17)$$

Finally, we need to ensure that τ_{min} is sufficiently small to provide good performance for large number of stations. To achieve this, m needs to be sufficiently large. By relying on a heuristic that hints efficient performance for up to 2^{m-1} stations when the total number of backoff stages is m , we choose the configuration $m = 6$: with this setting for m , the resulting τ_{min} is sufficiently small to ensure that $\tau = 1 - (1 - \gamma)^{1/(N-1)}$ crosses $G(\gamma)$ close to $\gamma = \gamma_{opt}$ even for N as large as 30.

Figure 4.3 shows the point of operation resulting from our configuration as well as the default ones for 1901 and 802.11, given by the intersection between the curves $1 - (1 - \gamma)^{1/(N-1)}$ and $G(\gamma)$. While the goal of our configuration is that $G(\gamma)$ decreases sharply at $\gamma = \gamma_{opt}$, we can observe that this decrease is smoothed by the randomness associated to the deferral and backoff counters. Despite this, our configuration is still much closer to γ_{opt} than the default configurations. Table 4.1 summarizes our design choices for the PLC CSMA/CA protocol.

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance

Parameter	Design choice	Value	Justification
d_i 's	$\left\lceil \gamma_{opt} \frac{2^i CW_{min}}{2} - 1 \right\rceil$	{0, 1, 3, 6, 13, 27}	Proper level of reaction to congestion; $G(\gamma)$ sharply decreasing at γ_{opt}
CW_{min}	$\left\lfloor \frac{2}{\gamma_{opt}} - 1 \right\rfloor$	12	τ_{max} large enough for $N = 2$
m	Efficient performance for up to 2^{m-1} stations	6	τ_{min} small enough for good performance with large N

Table 4.1 – Our proposed design choices for the PLC CSMA/CA protocol and best-effort applications (CA0/CA1 classes). The optimal collision probability γ_{opt} is computed given $T_c/\sigma \approx 81$. The IEEE 1901 default values of these parameters are given in Table 2.2.

4.3.2 1901-Proposed Advantage over 802.11

As we argued above, the key to optimize performance is to force that $\mathbb{E}[R]/\mathbb{E}[X]$ decreases from a large value to a small one at $\gamma = \gamma_{opt}$ as sharply as possible. In the following, we show that, thanks to the deferral counter mechanism, 1901 achieves a sharper transition, which is why it provides better performance to 802.11 across different values of N .

Equation (4.11) can be rewritten as

$$G(\gamma) = \frac{\mathbb{E}[R]}{\mathbb{E}[X]} = \frac{1}{\sum_{i=0}^{\infty} (1-\gamma) b c_{\min(i, m-1)} \prod_{j=0}^{i-1} p_{\min(j, m-1)}},$$

where p_j is the probability that, when the station moves out of backoff stage j , it jumps to $j+1$ instead of returning to backoff stage 0. As we discuss above, the main dependency of G on γ comes from the p_j 's. Hence, we analyze the behavior of p_j for both 1901 and 802.11, in order to understand the different performance of the two protocols. For 802.11, p_j is given by the probability that a transmission attempt collides, i.e.,

$$p_{802.11} = \gamma, \text{ for all } j.$$

For 1901, p_j is not only driven by the collision probability γ , but also by the deferral counter expiry probability $1 - t_j$, i.e.,

$$p_{1901} = 1 - (1 - \gamma)t_j(\gamma),$$

which, combined with (4.13), yields

$$p_{1901} = 1 - (1 - \gamma) \frac{d_j + 1}{\gamma CW_j}.$$

For $G(\gamma)$ to decrease sharply with γ at $\gamma = \gamma_{opt}$, we would like p_j to increase as sharply as possible. If we look at the derivative of p_j with respect to γ evaluated at $\gamma = \gamma_{opt}$ for 802.11 and 1901, we obtain, respectively

$$\begin{aligned} \left. \frac{dp_{802.11}}{d\gamma} \right|_{\gamma=\gamma_{opt}} &= 1, \\ \left. \frac{dp_{1901}}{d\gamma} \right|_{\gamma=\gamma_{opt}} &= \frac{1}{2\gamma_{opt}}. \end{aligned}$$

Given that $\gamma_{opt} \ll 1$, from the above we have $dp_{1901}/d\gamma \gg dp_{802.11}/d\gamma$, i.e., the probability of jumping to the next backoff stage increases much more sharply with 1901 than with 802.11, which explains the performance improvement achieved by 1901 over 802.11. This behavior is caused by the fact that while with 802.11 the parameter γ only affects the probability of jumping to the next backoff stage through the collision probability, with 1901 γ affects this probability not only through the collision probability but also through the probability that a station jumps to the next backoff stage upon expiring the deferral counter.

4.4 Performance Evaluation

In this section, we evaluate the accuracy of our model and the performance of our proposed 1901 configurations via simulation as well as by means of testbeds of 1901 and 802.11 devices. The section is structured as follows. First, we validate our simulator and model by a testbed of HomePlug AV (1901-compliant) devices. Second, we explore the accuracy of our model by using simulations over a wide range of configurations. Third, we extensively study our 1901 proposed configurations by (i) comparing the proposed configurations with the performance of an optimal 802.11 protocol and the default 1901, and by (ii) conducting an exhaustive search in the parameter space to confirm that our configuration performs closely to the one maximizing throughput across different numbers of stations. Fourth, we analyze the convergence time of our 1901 enhancements and compare it against the 802.11-optimal, dynamic algorithm of Patras et al. [34]. In contrast to this kind of adaptive algorithms, 1901 adapts the behavior based on the sensed transmissions and not on estimation techniques, which drastically reduces convergence time. Finally, we implement the 1901 CSMA/CA protocol on a testbed of WiFi hardware and show that our enhancement outperforms the default 1901 configuration in practice.

4.4.1 Experimental Validation of our Simulator and Model

We use simulations to evaluate 1901 performance. To this end, we wrote a `Matlab` simulator that implements the full CSMA/CA mechanism of 1901. In this subsection, we validate the accuracy of our simulator with experimental results from a HomePlug AV testbed. The guidelines to reproduce the testbed experiments are available in Appendix B.

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance

We build a testbed of 8 stations in ideal channel conditions, each comprising a PLC interface. In our experiments, N stations send UDP traffic (at a rate higher than the link capacities) to the same non-transmitting station using `iperf` for 240s. We run experiments for $1 \leq N \leq 7$, and we evaluate the collision probability as described in Appendix B.3. Figure 4.4 shows the average collision probabilities, obtained from 10 testbed experiments and 10 simulation runs (note that confidence intervals are so small that they can barely be appreciated and are omitted). We observe an excellent fit between experimental and simulation results.

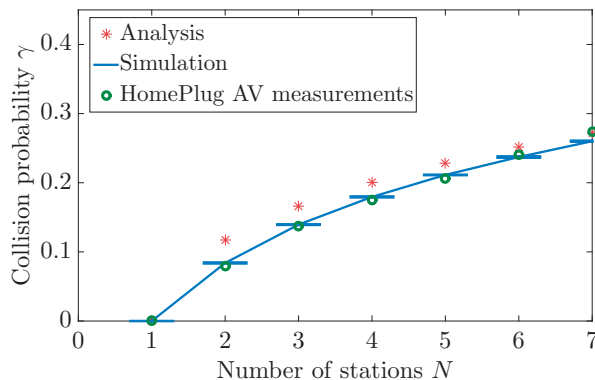


Figure 4.4 – Collision probability obtained by simulation, model, and experiments with HomePlug AV devices for the default class CA1 of 1901 given in Table 2.2.

Contrary to some existing 802.11 interfaces, the MAC parameters of the HomePlug AV devices cannot be modified, as they are stored in the firmware, and the required offsets of their binary values are not publicly available. Thus, the following results have been obtained with our validated simulator.

4.4.2 Simulation Assumptions and Parameters

For the simulations, we follow the network assumptions of Section 4.2.1 and take the time slot duration and timing parameters specified by the IEEE 1901 standard. A PLC frame transmission has a duration D , which is set to the maximum defined by IEEE 1901 since all stations are saturated, and is preceded by two priority tone slots (PRS) and a preamble (P). It is followed by a response inter-frame space ($RIFS$), the ACK, and finally, the contention inter-frame space ($CIFS$). Thus, a successful transmission has a duration $T_s = 2PRS + P + D + RIFS + ACK + CIFS$. In case of a collision, the stations defer their transmission for $EIFS \mu s$, where $EIFS$ is the extended inter-frame space used by 1901; hence, a collision has a duration $T_c = EIFS$. Table 4.2 shows all the timing parameters we use to evaluate throughput given (4.6).

Parameter	Duration (μs)
Slot σ , Priority slot PRS	35.84
$CIFS$	100.00
$RIFS$	140.00
Preamble P , ACK	110.48
Frame duration D	2500.00
$EIFS$	2920.64

Table 4.2 – Simulation parameters.

4.4.3 Validation of the Decoupling-Assumption Model

We now present the normalized throughput S obtained by our model and via simulation. Figure 4.5 compares the normalized throughput of 1901 obtained via simulation against our model, for (i) the default parameters for the two priority classes CA1 and CA3 (CA0 and CA2 are equivalent), and (ii) some additional configurations. We observe a good fit between analysis and simulation, with an exception for small N and small d_i 's (CA0/CA1 class and $CW_{min} = 8$, $d_0 = 0$, $m = 4$ configuration). The fact that the accuracy is somewhat reduced for small N in these cases is due to the decoupling assumption, which fails to capture the coupling introduced by the deferral counter. We thoroughly discuss this issue in Chapter 5 and, to improve this accuracy, we introduce another model that does not rely on the decoupling assumption in Chapter 6. The decoupling-assumption model is simpler than the one of Chapter 6, hence it enables us to propose enhancements that we evaluate next.

4.4.4 Proposed Enhancement for the IEEE 1901 MAC

We evaluate the enhancements proposed in Section 4.3. Given the timing parameters of Table 4.2 and the expressions (4.16) and (4.17), the parameters of the enhanced 1901 are $CW_{min} = 12$ and $d_i = \{0, 1, 3, 6, 13, 27\}$ for each backoff stage, with $m = 6$ (as mentioned earlier, our enhancements consist simply in modifying the protocol parameters).

Proposed vs. Default 1901

Figure 4.6 compares the performance of (i) 1901 using default CA1 configuration, (ii) 1901 using our proposed enhanced configurations mentioned above, and (iii) “optimal” 802.11 for a varying number of stations—for the “optimal” 802.11, CW_{min} is computed from (4.7) and the fixed-point equation for the collision probability [15, 34]. It appears that 1901 with our proposed configuration performs similarly—or better than—the optimally configured 802.11, which requires knowing the number of stations N . Furthermore, the proposed configurations drastically boost the efficiency of 1901.

In Figure 4.6, we also show similar results, but using $m = 4$ instead of $m = 6$ for the number of backoff stages. Here too, our proposed configuration yields substantial improvements. These improvements can offer tens of Mbps of throughput gain, given the data rate (up to 1.5 Gbps) of the latest PLC specification, HomePlug AV2.

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance




Figure 4.5 – Normalized throughput obtained by simulation and with our model, for the default configurations of 1901 (top) and for various configurations with $CW_i = 2^i CW_{min}$, and $d_i = 2^i(d_0 + 1) - 1$, $i \in \{0, m - 1\}$ (bottom). Lines represent simulations and points the analytical results.




Figure 4.6 – Simulations of 1901, enhanced 1901 proposed here, and optimal 802.11 for $m = 6$ (left) and $m = 4$ (right).

Exhaustive Search of Optimal Configuration

To further assess the performance of our proposed configuration, we run an exhaustive search in the parameter space of the CW_{min} and d_i 's; in particular, we take $m = 6$, $CW_i = 2^i CW_{min}$, and $d_i = \lceil f^i - 1 \rceil$, $i \in \{0, 5\}$, where $f \in \mathbb{R}$, and run the exhaustive search for $2 \leq CW_{min} \leq 32$ and $1 \leq f \leq 5$.

Our aim is to maximize throughput. However, since different numbers of stations have a maximum throughput at different configurations, we need a criterion to find a good tradeoff between the performance for different numbers of stations. To this end, we define function $\sum_{n=1}^N \log(S_n)$, where S_n is the normalized throughput of a scenario with n stations, and select the configuration that maximizes this function (computed with our model, for $N = 30$)⁶.

To validate our proposal, we compare it against the configuration obtained by performing an exhaustive search satisfying the above criterion. The values returned by this exhaustive search are $CW_{min} = 14$ and $f = 1.5$. Figure 4.7 presents the normalized throughput of the default configuration, the proposed one and the exhaustive search. We observe that our proposed configuration performs very close to the one returned by our exhaustive search algorithm; the configuration resulting from the exhaustive search slightly outperforms ours. We attribute this to the approximations we relied on in order to derive the proposed 1901 configurations.




Figure 4.7 – Performance comparison of the proposed 1901 configuration, the configuration that maximizes throughput obtained from an exhaustive search, and the default 1901 configuration.

⁶To solve the maximization problem over the integer-parameter space, we use the `ga` toolbox of Matlab.

Multiple Contention Domains

The performance of the 1901 MAC under multiple contention domain scenarios is quite different from the results presented in previous subsections. 1901 specifies an optional RTS/CTS mechanism, which is not enabled by default on our hardware. Thus, without this mechanism, the performance degrades in scenarios such as hidden nodes. Our proposed configuration does not solve such performance issues, but it attenuates their effects as it uses larger a contention window than the default one; our $CW_{min} = 12$ value is larger than the one of default 1901 ($CW_{min} = 8$), hence hidden terminals can transmit more packets on average while the contending station is in backoff⁷. Note that the deferral counter is not triggered if two contending stations are hidden to each other.

Saturated and Non-saturated Stations

We now evaluate the performance of our proposed configurations in a mixed scenario of both saturated and non-saturated stations. The non-saturated stations have Poisson traffic with arrival rate λ . Figure 4.8 presents the normalized throughput for a scenario where half of the stations are saturated and half are not. Also, it depicts the same metric under a scenario where only 3 stations are saturated, whereas the number of non-saturated stations varies. We observe that under these scenarios, our proposed configurations perform better than the optimal 802.11 and the default 1901.




Figure 4.8 – Simulations of 1901, enhanced 1901 proposed here, and optimal 802.11 for $m = 6$, in scenarios that include non-saturated stations.

Evaluation under Dynamic Traffic

The evaluation of the proposed configuration in the previous sections has been limited to static scenarios with a fixed number of stations. However, in reality the number of active stations is not fixed, but it varies with time. One of the important advantages of the proposed enhancement as compared to other optimal adaptive approaches is that 1901 adapts the behavior with a very fine granularity: it updates CW after sensing

⁷The average time in backoff under no contention is $(CW_{min} - 1)/2$.

transmissions, which yields fast adaptation to varying traffic. In contrast, adaptive algorithms for 802.11 are required to update their estimation periodically after having enough samples of collided/transmitted frames, which can yield long convergence times under dynamic traffic.

In the following, we evaluate the convergence time of our approach and compare it against one representative adaptive algorithm for 802.11: the distributed adaptive control (DAC) algorithm of [34], which estimates the collision probability and drives the system to its optimal value by adjusting CW_{min} periodically (every 100 ms). For this evaluation, we run simulations in which we drastically change the number of saturated stations from 10 to 2. Figure 4.9 shows the average CW (averaged over all stations and over 200 ms intervals) and the instantaneous throughput (averaged over 200 ms intervals) obtained from these simulations. Results confirm that existing 802.11 adaptive algorithms fail to react quickly to changing scenarios, harming the resulting performance over transients, and in contrast, our approach reacts much faster.




Figure 4.9 – Performance comparison of enhanced 1901 proposed here and of the DAC 802.11 algorithm under dynamic traffic. At 16 s (shown by the vertical line), the number of saturated stations changes drastically from 10 to 2. The dashed rectangle shows the convergence time of the DAC algorithm, which is in order of seconds.

4.4.5 Experimental Validation of Proposed Enhancement

To confirm the improvements shown via simulation in the previous sections, we implement the 1901 protocol with the proposed configuration in WiFi hardware. The reason for choosing this platform for our experimental validation is that, to implement our enhanced configuration of 1901, we need full access to the firmware, which commercial 1901 devices do not provide. Therefore, we employ 802.11 Broadcom wireless cards and substitute the default proprietary firmware with *OpenFWWF* [37] that has already been used to extend and modify the 802.11 default behavior [38, 39].

The assembly code of the firmware has been modified to follow the 1901 protocol. To this end, we store the different values of contention window (CW_i) and deferral counter

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance




Figure 4.10 – Deployed testbed, consisting of 1 AP (7 dBi omnidirectional antenna) and 25 stations (2 dBi omnidirectional antenna).

Parameter	Description	Value
R_d	Data rate	5.5 Mb/s
R_c	Control rate	2 Mb/s
L	Packet length	1460 B
σ	Slot duration	20 μ s
T_{plcp}	Preamble duration	192 μ s
T_c	Collision duration	2720 μ s
-	Band	2.4 GHz
-	Transmission power	20 dBm

Table 4.3 – WiFi Testbed parameters.

(d_i) into the shared memory and tweak the `rx_plcp` and `update_contention_params` modules to implement the backoff procedure defined in 1901 after a data transmission or when an ongoing data transmission is sensed in the channel.

With the above implementation, we deploy our testbed of Alix 2d2 devices from PC Engines, each with a Broadcom BCM94318MPG 802.11b/g miniPCI wireless card and with the Ubuntu 10.04 Linux (kernel 2.6.36) distribution installed. The testbed is located under a raised floor of a laboratory and it comprises 25 wireless stations, and one desktop machine that acts as an access point (AP), as shown in Figure 4.10. We carefully run our experiments when the channel activity is very low to avoid any external source of interference. To generate traffic and evaluate performance, we use `mgen` [40] configured to send uplink saturated UDP traffic.




Figure 4.11 – Experimental results of enhanced 1901 proposed here for $m = 6$ and default 1901 implemented on WiFi stations.

To emulate the PLC PHY/MAC timing parameters, we set the slot duration and fix the data-rate of WiFi to a small value, such that the slot duration is $20 \mu s$ and the duration of a successful transmission or collision is $2720 \mu s$, which are very close to the ones of the IEEE 1901 standard (these and the rest of the transmission parameters are given in Table 4.3). The parameters of the enhanced 1901 in this case are $CW_{min} = 16$ and $d_i = \{0, 1, 3, 7, 14, 29\}$ for each backoff stage, with $m = 6$. For each experiment, we perform 10 runs. For all results, 95% confidence intervals are below 1%. The results of this experiment, given in Figure 4.11, are qualitatively very similar to the simulation results shown earlier and confirm that (i) the proposed configuration maintains very good throughput performance as N increases, (ii) it substantially outperforms the default 1901 configuration, and (iii) it performs very closely to the optimal 802.11 configuration.

4.5 Related Work

Models of 802.11 can form a basis for analyzing 1901, as 802.11 can be viewed as a simple version of 1901 in which the deferral counter never expires. In this chapter, we adopt an analytical framework from 802.11 performance evaluation in order to model 1901. Therefore, review related work on models and enhancements of both 802.11 and 1901.

Models of IEEE 802.11: A large number of performance evaluation models have been proposed for 802.11 (e.g., [15, 30]). Among those, the model proposed by Bianchi [15] for single contention domain networks is very popular. This work models the backoff process of 802.11 by using a discrete time Markov chain. The main assumption behind this model is that the backoff processes of the stations are independent, which is known as the *decoupling assumption*. With this assumption, the collision probability γ experienced by all stations is time invariant, and can be found by solving a fixed-point equation that depends on the protocol parameters.

Kumar et al. [30] study the backoff process of 802.11 under the decoupling assumption using renewal theory. In this chapter, we employ a similar method for finding a fixed-point equation for 1901. But due to the increased complexity of 1901, this is significantly more challenging than that for 802.11.

Models of IEEE 1901: There are two works that analyze the backoff process of 1901 [36, 41]. Chung et al. [36] propose a model similar to Bianchi’s model for 802.11 [15]. However, they do not provide the corresponding fixed-point equation for the collision probability, due to the increased complexity of the Markov chain resulting from introducing the deferral counter. To compute the collision probability in this case, a costly system of more than a thousand non-linear equations has to be solved. Moreover, it has not been investigated whether this system of equations has a unique solution. The work of Cano and Malone [41] discusses the validity of the decoupling assumption and simplifies the model of Chung et al. [36] under non-saturated assumptions. Compared to

Chapter 4. Decoupled Analysis and Enhancement of Multi-user Performance

our work, these works do not provide any insights for the uniqueness of the solution or enhancements.

We propose a model that is strictly equivalent to the model of Chung et al. [36] in terms of accuracy. The key difference is that, in our case, the collision probability can be obtained by solving a single fixed-point equation. In this sense, our model can be seen as a drastic simplification of the model by Chung et al. [36], and this simple form enables us to derive efficient configuration parameters for 1901.

Enhancements of IEEE 802.11: There is a large body of work introducing enhancements for the 802.11 CSMA/CA. In particular, Bianchi [15] computes the optimal contention-window that achieves maximum throughput, which is a function of the number of contending stations N . Typically, N is unknown and varies with time, hence practical implementations of such optimal configurations use some estimation techniques.

To apply Bianchi's analysis, several attempts have been made to estimate the number of contending stations [34, 35, 42]. These methods typically rely on measuring the collision probability or the channel activity and on estimating N periodically. The main disadvantage of such approaches is that they introduce more complexity at the MAC layer, thus challenging their practical implementation. Moreover, such approaches can lead to long convergence-times under dynamic traffic.

Enhancements of IEEE 1901: There are not many studies on enhancements of 1901. Campista et al. [43] propose a constant d_i equal to 0 for all backoff stages i , which means that whenever a station senses the medium busy it doubles CW . This technique decreases the collision probability, but yields the most extreme case of 1901 unfairness; for small N , always doubling CW causes idle stations to have fewer chances of accessing the channel, compared to a station that just transmitted successfully (and whose CW is minimal). This yields a high variance of delay, as we explain later in Chapter 5. Kriminger and Latchman [44] provide a mechanism that keeps d_i and CW constant, but where CW depends on the number of stations N . Requiring that N be known is impractical.

Additional MAC-Layer Studies of IEEE 1901: We point out studies that go beyond the CSMA/CA mechanism of 1901 but that still explore MAC-layer performance. Cano and Malone [45, 46] investigate impairments in the priority-resolution process (described in Section 2.5.2). The authors find a starvation and variability of low-priority traffic when higher priorities exist in the network, and an oscillatory behavior of higher-priority classes. Zarikoff and Malone [47] establish guidelines for PLC testbed construction and experiments, under ideal channel settings. They perform UDP and TCP throughput evaluations under different contending-flow scenarios and priority classes. They also [48] study via simulations the performance of TDMA reservations of the 1901 MAC layer, a feature that it is not currently being used in practice.

4.6 Summary

The MAC layer of IEEE 802.11 can react to contention with a fine granularity and without involving collisions, which offers the potential of high gains in terms of throughput as compared to 802.11. Unfortunately, with the default parameter setting of 802.11, the protocol operates far from optimality and does not fully exploit its potential. One possible reason for this is the lack of a simple model of 802.11 that can be leveraged to find an appropriate setting of its parameters. In this chapter, we introduce a simple model of 802.11 and we prove the uniqueness of the solution for a wide range of configurations.

One of the main challenges to modeling 802.11 is the complexity of the protocol that has a very large state-space. To reduce the state space and to come up with a simple model, we make the assumption that the backoff processes of the stations are independent. Building on the resulting model, we derive a procedure for steering the network towards its optimal point of operation. With this, we obtain a protocol that provides performance close to the optimum, independently of the number of stations, and reacts quickly to traffic-load changes. This is a result that has been long pursued by the research community.

Our proposal requires only modifying existing parameters, and does not change the CSMA/CA 802.11 algorithm itself. Therefore, it can be easily incorporated into practical deployments. We implement the 802.11 protocol on a testbed of WiFi hardware. Our simulations and testbed measurements confirm the drastic performance improvements of our proposal. This chapter focuses on elastic applications that are quite insensitive to jitter; their performance mainly depends on throughput. In the next chapter, we discuss the short-term dynamics of 802.11 and jitter, hence treat delay-sensitive applications.

5 Short-Term Fairness and Dynamics in Multi-user Settings

5.1 Introduction

In the previous chapter, we extensively investigated the performance of the IEEE 1901 CSMA/CA protocol. We focused on the average performance at steady-state, omitting evaluations on short timescales. We observed that this CSMA/CA protocol can be very efficient. In this chapter, we delve into the short-term dynamics of multi-user performance and we consider delay-sensitive applications. Notably, we find that the *high efficiency of 1901 comes at a cost of short-term unfairness that introduces high jitter*.

Fairness is achieved when the medium is shared equitably among the stations during a fixed time-interval. Depending on the length of this interval, fairness is defined as long term or as short term, and it determines the quality of service. Long-term fairness is related to the average throughput per user, and short-term fairness affects delay variance (i.e., jitter). In a short-term unfair system, a station has to wait for many other stations to transmit before it transmits again, which increases delay. We examine short-term fairness for two reasons. First, short-term fairness is a property stronger than long-term fairness: short-term fairness implies long-term fairness, but the reverse does not hold. Second, in the long term, by the law of large numbers, all stations share equitably the medium (when they all have the same channel conditions, and there is no hidden terminal).

Short-term dynamics and fairness are crucial not only for delay, but also for modeling-performance assumptions. In the previous chapter, we assumed that stations collide or sense the medium busy with a time-invariant probability. In the case the protocol is short-term unfair, clearly this assumption is not viable: any station that defers transmissions due to sensing the medium busy, observes a busy probability that depends on the system state. In this chapter, we observe that indeed, PLC unfairness introduces “coupling” between the stations, which questions independence assumptions.

We present a snapshot of the short-term dynamics of 802.11 and 1901 from a testbed experiment. Figure 5.1 provides evidence of the aforementioned unfairness, jitter and coupling phenomena, for a testbed with two stations. Compared to 802.11, 1901 is




Figure 5.1 – Testbed trace of 50 successful transmissions by two saturated stations with 1901 and 802.11a. 1901 exhibits short-term unfairness: a station holding the channel is likely to continue holding it for many consecutive transmissions, which causes high jitter. In addition, 1901 unfairness makes the decoupling assumption questionable. 802.11 is fairer, which makes the decoupling assumption viable in this case.

short-term unfair. While the 1901 Station *A* transmits during several consecutive slots, Station *B* is likely to remain in a state where it has a high probability of colliding or sensing the medium busy. *B* is then even less likely to attempt a transmission while in this state, and it might have to wait several tens of milliseconds before the situation reverts. Thus, the collision probabilities observed by the stations are clearly time varying, which invalidates the decoupling assumption. Furthermore, this unfairness translates into high delay-variance (i.e., high jitter).

The fairness issue with 1901 stems from the introduction of the deferral counter. Due to this mechanism, the stations might double the contention window after sensing the medium busy. In networks of a few stations, this causes the last station that transmitted to have a higher probability of transmitting again, compared to the other stations. Any station that transmits, sets the contention window to the minimum one; whereas the rest of the stations might double their contention window. As we observed in Chapter 4, the deferral counter can adjust the network performance by configuring its initial values at each backoff stage. Here too, the deferral counter can yield different fairness levels depending on its initial values. In fact, we show that the deferral counter determines a tradeoff between throughput and fairness.

The remainder of this chapter is organized as follows. In Section 5.2, we present our metrics with which we evaluate fairness. In Section 5.3, we prove analytically that PLC is short-term unfair for a network with two stations and that WiFi is much fairer

than PLC on this timescale. Section 5.4 includes our simulation and testbed results for various number of stations and for different fairness-horizons. We show that short-term fairness introduces substantial delay and jitter in Section 5.5, and that 1901 involves a tradeoff between throughput and fairness in Section 5.6. In Section 5.7, we find that the PLC unfairness penalizes the accuracy of the decoupling assumption for a small number of stations. We review related work in Section 5.8 and give concluding remarks in Section 5.9.

5.2 Metrics of Fairness

Short-term fairness of 802.11 has been extensively investigated (e.g., [49, 50]) in single collision-domain networks. We analyze similarly the 1901 protocol and evaluate fairness by using two metrics: *Jain's fairness index* and the *number of inter-transmissions*.

Jain's fairness index [51] is used to measure fairness, because it represents the variability of a set of measurements (e.g., the throughput of the stations). Larger values of Jain's index J indicate better fairness. Ideally, when all stations share equitably the channel, this index is equal to 1. Note that $1/N \leq J \leq 1$. Jain's index is defined as follows.

Definition 2. *Let N be the number of contending stations in the network and ρ_n be the throughput of station n within a time period. Then, Jain's fairness index is defined as*

$$J = \frac{\left(\sum_{n=1}^N \rho_n \right)^2}{N \sum_{n=1}^N \rho_n^2}. \quad (5.1)$$

Fairness can be estimated at different time-horizons. As the time horizon increases, fairness should improve in a CSMA/CA protocol with homogeneous users, as according to the law of large numbers, every user gets an equal share of the channel in the long term. To express fairness as a function of the time horizon, we use the sliding-window method (SWM) [52]. This method slides a window of size w packets across a trace of successful transmissions. For each sequence of w packets, Jain's fairness index is computed from (5.1), with ρ_n being the fraction of transmissions performed by station n within the window of w packets. Jain's fairness index $J(w)$, associated with the window size w , is the average of the fairness values of all consecutive sequences of w packets in the trace. As w increases, the fairness time-horizon increases. Figure 5.2 shows an example of the SWM for two window values. In our study, we use a normalized window W with respect to N , given by $W = w/N$. In Section 5.4, we apply the SWM to traces of successfully received packets acquired from simulation and real-network tests.

Although Jain's fairness index with the SWM is a good metric of fairness on different time-scales, it is analytically intractable. Therefore, to model short-term fairness, we

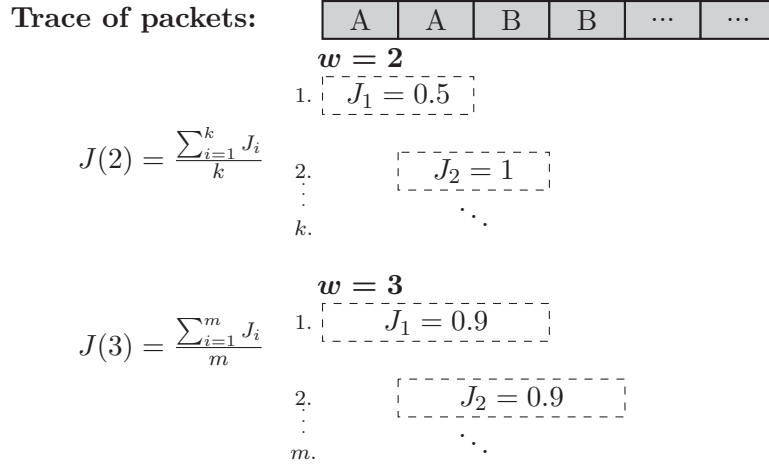


Figure 5.2 – Sliding-window method example with two contending stations A and B. The labels on the successfully transmitted packets represent the source of the packet. We show how $J(w)$ is computed for two window values, i.e., $w = 2$ and $w = 3$.

employ another metric, called the number of inter-transmissions.

The number of inter-transmissions was introduced by Berger-Sabbatel et al. [49] for studying 802.11. The objective of this metric is to analytically evaluate short-term fairness. It is defined as follows.

Definition 3. *The number of inter-transmissions K for a tagged station B is the number of transmissions originating from other stations between two consecutive transmissions of B .*

In other words, if $K = 0$, after a successful transmission of Station B , the next transmission is done again by B . If $K = 1$, Station B transmits once, then another station transmits, and then B transmits again. If there are N stations contending for the channel, the ideal value of K is $N - 1$. In reality, K is a random variable, and $\mathbb{E}[K] = N - 1$ represents an ideally fair protocol. The second moment of K is related to jitter: the larger $\text{Var}(K)$ is, the higher the delay variance is. The analysis of the distribution of K for networks with two stations is presented in the next section.

5.3 Analysis

We analyze 1901 short-term unfairness with the number of inter-transmissions. We focus on the case with $N = 2$, because in this case we can maintain the uniform distribution of the backoff counters, followed by the standard. Hence, we achieve a high level of accuracy with our analysis. For $N > 2$, due to the challenging nature of the problem, we need to assume exponential backoff counters per station¹. Even with this approximation, the

¹In this case, if the mean of the backoff counter at station i is $1/\lambda_i$, then station i wins the channel with probability $\lambda_i / \sum_{n=1}^N \lambda_n$, because of the distribution of the minimum of N exponential random

state-space explodes: The first terms of the probability mass function of the distribution, $\mathbb{P}(K = k)$, can be computed, but the model is intractable for large k due to the explosion of the number of states of the system.

5.3.1 Assumptions

We assume the following in terms of network and traffic conditions:

Perfect sensing: There is a single collision-domain with two stations A and B , and there are no hidden terminals.

Initial system state: Both stations start at backoff stage 0. From a practical point of view, we analyze the system under scenarios where Station A is saturated and Station B has low-intensity traffic (e.g., very low packet-arrival rate). We compute the distribution of K for Station B .

Perfect channel: There is no packet loss or errors due to the physical layer.

Infinite retry limit: The stations have an infinite retry limit; that is, they never discard a packet until it is successfully transmitted.

In terms of modeling hypotheses, we assume that the backoff counters of the stations are *continuous* random variables uniformly distributed over the range $[0, CW]$. Hence, the probability of collision is equal to 0. This assumption is necessary because the analytical model of the real discrete-time protocol appears intractable.

The backoff counters of the two stations are independent: This is not an assumption but results from the real protocol where stations select independently their backoff counters. The computation of the distribution $\mathbb{P}(K = k)$ is based on the observation that, whenever the stations choose or resume their backoff counters, the station with the smallest backoff counter transmits. We are now ready to introduce our model.

5.3.2 The Distribution of Inter-transmissions

Let A and B be the two contending stations. We tag station B . Let b be the backoff counter of B and a_i be the backoff counter of the A before the i th transmission. Then, the distribution of K for 802.11 is given by

$$\mathbb{P}(K = k) = \mathbb{P}\left(\sum_{i=1}^k a_i < b \text{ and } \sum_{i=1}^{k+1} a_i \geq b\right).$$

For 802.11, the distribution of K is computed by Berger-Sabbatel et al. [49] and reads

$$\mathbb{P}_{802.11}(K = k) = \frac{k+1}{(k+2)!}, \quad k \geq 0. \quad (5.2)$$

In contrast to 802.11, the computation of this distribution for 1901 is more complex.

variables.

Let BC_B , CW_B and DC_B be the values of backoff counter BC , contention window CW and deferral counter DC at B , respectively. We define similarly the counters of A . In 802.11, under the collision-free assumption, the contention window remains constant for both stations, independently of the value of K . Whereas, in 1901 CW_B is doubled whenever A transmits and $DC_B = 0$, because B senses the medium busy. To evaluate $\mathbb{P}_{1901}(K = k)$, we compute the probability that A transmits k times consecutively, or equivalently the probability that $BC_B < BC_A$ at each of the k transmission attempts. BC_A is always uniformly distributed over the range $[0, 8]$, whereas BC_B depends on k . When $K = 0$, we have $CW_B = 8$ and $DC_B = 0$. After each transmission of A , BC_B is updated as follows: if DC_B is equal to 0, then new BC_B and DC_B are chosen depending on the backoff stage of station B ; otherwise, BC_B is decremented by BC_A (Station B has already counted down BC_A idle slots) and DC_B is decremented by 1. By using this reasoning and Table 2.2 with the 1901 parameters, we compute DC_B and the range to which BC_B belongs for each k .

The 1901 distribution of K for the CA1 priority-class is given as follows. The proof can be found in Appendix A.3.

Proposition 1. *The distribution of the number of inter-transmissions K for the CA1 priority-class of IEEE 1901 standard is*

$$\mathbb{P}_{1901}(K = k) = \begin{cases} 1/2 & k = 0 \\ 1/8 & k = 1, 2 \\ 1/32 & 3 \leq k \leq 6 \\ 0.0578^l / 128 & 7 \leq k - 16l \leq 14 \\ 0.0578^l \cdot \mathbf{I}(k - 16l - 14) / 8 & 15 \leq k - 16l \leq 22, \end{cases} \quad (5.3)$$

where $l = \lfloor (k - 7) / 16 \rfloor$, and $\mathbf{I}(n)$ is the n th entry ($1 \leq n \leq 8$) of the vector $\mathbf{I} = 10^{-2}(6.25 \ 6.25 \ 6.24 \ 6.17 \ 5.92 \ 5.38 \ 4.53 \ 3.49)$.

An important consequence of Proposition 1 is that the distribution of K has a tail for 1901 heavier than for 802.11, meaning that with 1901 a station is more likely to perform many more consecutive back-to-back transmissions than with 802.11. To see this, consider the survival functions of the 1901 and 802.11 distributions of K , $S_{1901}(k) = \mathbb{P}_{1901}(K > k)$ and $S_{802.11}(k) = \mathbb{P}_{802.11}(K > k)$. We can compute from (5.2) and (5.3) that

$$\lim_{k \rightarrow \infty} \frac{S_{1901}(k)}{S_{802.11}(k)} = \infty,$$

which means that the distribution of K has a tail for 1901 heavier than for 802.11 [53]. This is also observed from the two distributions plotted in Figure 5.3, and from the moments of the distributions: We have $\mathbb{E}[K] = 2.8$ and $\text{Var}(K) = 28.28$ for 1901, whereas 802.11 yields $\mathbb{E}[K] = 0.73$ and $\text{Var}(K) = 0.77$. It is important to note that the variance explodes with 1901, which shows that, compared to 802.11, 1901 generates much higher jitter.




Figure 5.3 – The distribution of K for 1901 (CA1 priority) and 802.11 computed analytically in log-log scale. The x-axis is the same for both plots. Note the difference of y-axis scale between the two standards, e.g., $\mathbb{P}_{1901}(K=15) \sim 10^{-3}$ and $\mathbb{P}_{802.11}(K=15) \sim 10^{-14}$. The 1901 distribution of K attenuates periodically, as observed from (5.3).




Figure 5.4 – The distribution of K for 1901 (CA1 priority) in log-log scale computed analytically and via simulation with the real discrete backoff counters BC .

The distribution of K is the same for any minimum contention window CW_0 , as long as the deferral-counter values are maintained at the same value and the contention window is doubled between successive backoff stages. To study the real 1901 protocol with discrete backoff counters, we run our simulations with increasing values of CW_0 , multiplying the CW_0 by 2 and maintaining the deferral-counter values the same as in Table 2.2. For $CW_0 = 8$ (CA1 class of 1901), we find that $\mathbb{E}[K] = 4.2$ and $\text{Var}(K) = 70.1$ from simulations of 1901. Figure 5.4 depicts $\mathbb{P}_{1901}(K = k)$ obtained from analysis and simulations of the real protocol with $CW_0 = 8$. We observe that there is a slight deviation in $\mathbb{P}_{1901}(K = 2)$, which translates into a deviation of $\mathbb{E}[K]$ between analysis and simulation, due to the long-tailed distribution. As $CW_0 \rightarrow \infty$, the probability of collision tends to 0, and the discrete distribution of the backoff counters (real 1901 protocol) can be approximated by the continuous one (our model). For instance, the distribution of K from the simulation of the real 1901 protocol with $CW_0 = 1024$ coincides with the analytical distribution; We verify this by using the Kullback–Leibler divergence with the two distributions. Therefore, our model is a good approximation of the real protocol, and collisions deteriorate short-term fairness for $N = 2$: We have $\mathbb{E}[K] = 4.2$ with $CW_0 = 8$ in simulation of the real protocol, compared to $\mathbb{E}[K] = 2.8$ in the collision-free analysis.

This concludes our study of inter-transmissions. In the remainder of this chapter, we use Jain’s fairness index to evaluate fairness for $N \geq 2$. We have extended our study of inter-transmissions for $N > 2$, but the intractable nature of the problem prevents us from deriving a closed-form expression for $\mathbb{P}_{1901}(K = k)$.

5.4 Fairness Evaluation via Simulation and Testbed

We explore via simulations and experiments both the short-term and long-term fairness of 1901 and 802.11. To this end, we use Jain’s fairness index with the sliding-window method, as described in Section 5.2.

Let us first briefly explain our simulation and testbed setup. To compare the CSMA/CA mechanisms (802.11a and 1901), we develop a `Matlab` simulator, similar to the one of Section 4.4. The simulator is time-slotted, it features the backoff procedures of the 1901 and 802.11 standards and assumes perfect sensing and error-free channel. We use the simulator to obtain a trace of successfully transmitted frames over the channel. We do not simulate the PHY layer, because the two channels are different, and because we want to avoid hidden factors that might affect our CSMA/CA evaluation. All stations are saturated. The accuracy of our simulator is confirmed by the experimental results.

In addition to our simulator, we build a testbed of 6 stations, each comprising both power-line and wireless communications interfaces. The testbed experiments presented in this chapter are carried out on different mediums (i.e., power-line and wireless), and *under ideal channel conditions*. The stations are equipped with a Homeplug AV interface, and a WiFi one that is set to 802.11a mode and to a channel that does not interfere with other networks in our building. The detailed description of our experimental framework

can be found in Appendix B.3.

In our setup, N stations send UDP traffic at a rate higher than the link capacities to the $(N + 1)$ th station using `iperf`. We run tests for $N = 2, 3, 4, 5$. The $(N + 1)$ th station captures the frames received during a fixed time interval of 2 minutes. During this time interval, approximately $8 \cdot 10^4$ frames are captured. Our findings are developed in the following paragraphs.

The default class CA1 of 1901 is short-term unfair. Figure 5.5a shows the results of the SWM applied to traces of successfully transmitted packets from simulation and experiments for $N = 2$. We observe that Jain’s fairness index $J(W)$ for 802.11a is higher than for 1901 at all values of the normalized window W . Thus, 802.11 is fairer for all time horizons, and particularly in the short-term. For instance, $J(W) = 0.95$ is achieved with $W = 5$ for 802.11a and with $W = 70$ for 1901.




Figure 5.5 – Jain’s fairness index with the SWM. Simulations coincide well with testbed results. For 802.11a, simulation and testbed results are exactly the same, making the two curves indistinguishable.

The worst unfairness for 1901 occurs in case $N = 2$. The short-term fairness of the protocols changes as N increases. Figures 5.5b and 5.6 show the value of $J(W)$ from simulation and testbed results for $N = 3$ and 5, respectively. The two curves for 1901 and 802.11a approach each other as N increases (compare Figures 5.5a, 5.5b, 5.6). The largest difference between the mechanisms occurs for $N = 2$, which is expected to be a common scenario in home networking. As N increases, there is a higher probability that some station steals the medium from the station already monopolizing it. Therefore, the worst situation of 1901 unfairness is for $N = 2$.




Figure 5.6 – Testbed results of Jain’s fairness index with the SWM for $N = 5$.




Figure 5.7 – $J(W = 1)$ as a function of N from simulation results. Short-term fairness improves for 1901 and deteriorates for 802.11 as N increases.

As N increases, 1901’s fairness improves, but 802.11’s fairness deteriorates. 1901 is short-term fairer than 802.11 for $N \geq 15$. Figure 5.7 plots $J(W = 1)$ as a function of N in simulation. $J(W = 1)$ for 802.11a decreases as N increases, which has already been observed in previous works [49, 50]. This is due to the collision probability that increases with N . In 802.11 for large N , after a collision, a station that transmitted successfully in the past is favored compared to the others, because there exists a subset of stations with a larger CW due to collisions. Figure 5.7 also uncovers that 1901 is fairer in the short-term for $N \geq 15$. Furthermore, it illustrates two regimes that characterize the 1901 behavior; for $N < 10$, $J(1)$ decreases, and for $N \geq 10$, $J(1)$ increases. When $N < 10$, $J(1)$ for 1901 decreases for the same reason as for 802.11, as explained above. When $N \geq 10$, $J(1)$ increases because under high collision-probability conditions, in 1901 there is not a subset of stations with low CW , which is contrary to 802.11: After a collision all stations (including those that transmitted successfully in the past) increase their CW due to sensing the medium busy. Moreover, the increasing behavior of $J(1)$ for 1901, when N is large, is attributed to the high collision-probability and to the inability




Figure 5.8 – Left: Illustration of our experiment. Right: The median RTT of 10^4 ping requests of B as a function of A 's `iperf` load, measured from our testbed. The error bars represent the 10% and 99% percentiles. Observe the difference in the scales of y-axes of 802.11 and 1901.

of the maximum CW ($CW_{max} = 64$) to accommodate collisions. Consequently, when N is large, the majority of the stations have $CW = CW_{max}$, hence they have the same probability of transmitting, which results in fairness.

5.5 Short-Term Fairness and Impact on Delay

Short-term unfairness increases delay in networks with a small number of stations. A real example of this phenomenon with $N = 2$ is presented in Figure 5.8, where a station A sends UDP traffic with a varying load to a station C by using `iperf`, while a station B sends `ping` requests of size 1400 bytes to C . Figure 5.8 shows the median round trip time (RTT) of 10^4 ping requests of B for both 1901 (top) and 802.11a (bottom). The maximum RTT was 215.7 ms for 1901 and 11.5 ms for 802.11 when A was saturated. The results reveal that in a scenario where A sends saturated traffic and B sends bursty traffic in the form of individual packets, B experiences delays larger than the maximum tolerable delays of delay-sensitive applications. The average and maximum RTT of B when A was saturated were 8.2 and 21 ms, respectively, for CA3 priority.

The difference between the test mentioned above and the saturated UDP traffic tests presented in the previous section is that in the `ping` tests both stations have their CW equal to CW_0 when they start contending for the channel. This depicts a scenario with a bandwidth-hungry application (Station A) and a low-demand and delay-sensitive one (Station B) that contend simultaneously for the medium. In Section 5.3, we analyzed the distribution of inter-transmissions of station B under this scenario, i.e., when A and B contend with their initial CW equal to CW_{min} . By using our model, we showed that 1901 is short-term unfair.

We next compare the two priority classes of 1901 standard in simulation when all stations are saturated. Figure 5.9 presents short-term fairness and jitter for CA1 and




Figure 5.9 – Simulations with the IEEE 1901 configurations for CA1 and CA3 priority classes. The better the fairness is, the lower the jitter is.

CA3 classes with configurations shown in Table 2.2. We find that the delay-sensitive class of the standard, CA3, yields better fairness and lower jitter than CA1. Recall that, from the results of Section 4.4, the throughput for CA3 class is lower than that of CA1. Therefore, there is a tradeoff that is already exploited in 1901 for different levels of quality-of-service. In the next section, we thoroughly investigate this tradeoff and explore the parameters that fine-tune different levels of throughput and short-term fairness.

5.6 The Tradeoff between Throughput and Short-Term Fairness

In the introductory example of Figure 5.1, we observe a bi-stability effect with two 1901 stations, where stations are likely to remain for long durations in states with large transmit (resp. deferral) probabilities. We explain that this effect is caused by the deferral counter that creates a coupling between the stations and penalizes the accuracy of models assuming decoupling. In Chapter 4, we find that this coupling is beneficial for throughput². Here, we show that different 1901 configurations determine different tradeoffs between fairness (or jitter) and throughput.

We now further investigate this phenomenon in terms of short-term fairness. A MAC protocol is short-term fair when the stations have similar transmission opportunities over short time-scales. Conversely, an unfair protocol gives an advantage to some stations over others, which in practice results in high delay-variance (jitter). To measure short-term fairness, we compute Jain’s fairness index over windows of N packet durations, $J(W = 1)$ (see Section 5.2 for details). In the following, we take the normalized $W = 1$, as this is the smallest value of W such that $J(W)$ can be equal to 1 (for a perfectly fair protocol). The reported results are obtained by the SWM, along the whole packet traces. All stations are saturated.

²Intuitively, this is easy to understand: without proper synchronization, having one station transmitting for long durations is more efficient than alternating transmissions.

5.6. The Tradeoff between Throughput and Short-Term Fairness




Figure 5.10 – Throughput and fairness ($J(W = 1)$) with $CW_{min} = 8$ and $m = 4$ and various values of d_0 and f . The initial value of the deferral counter at backoff stage i is given by $d_i = f^i(d_0 + 1) - 1$.

We explore the tradeoff for a wide range of configurations with $CW_i = 2^i CW_{min}$, and $d_i = f^i(d_0 + 1) - 1$, $i \in \{0, m - 1\}$. In Figure 5.10, we plot throughput and short-term fairness as a function of the initial values of the deferral counter (in terms of d_0 and f), for $N = 2$ and $N = 5$. Notably, a direct link can be made between throughput and fairness. When the initial deferral-counter values are small (corresponding to small values of d_0 and f), the stations are more likely to react upon sensing the medium busy, and thus proactively avoid collisions. Conversely, these configurations have the worst short-term fairness and thus cause higher jitter.

We look more in detail on the throughput/fairness tradeoff in Figures 5.11 and 5.12. Both figures show the throughput and fairness achieved on networks with a varying number of stations. Figure 5.11 shows throughput and fairness for various initial deferral counter values (in terms of d_0 and f values). Figure 5.12 considers these two metrics for different numbers of backoff stages m . Both figures show again a clear tradeoff between throughput and short-term fairness. Furthermore, this tradeoff can be tuned by adapting the parameters controlling the number of backoff stages and the initial sizes of the deferral counters. This possibility is a remarkable feature of 1901, enabled by the deferral counter. We summarize the effects of all parameters on throughput and fairness in Table 5.1. The table also contains the minimum contention-window CW_{min} . Note that, in general, as discussed in Section 5.3, CW_{min} does not affect fairness as long as the contention windows are doubled between successive backoff stages. However, CW_{min} affects throughput: For small N , smaller CW_{min} values are better because of smaller backoff overhead. For large N , larger CW_{min} values are better for lower collision probability. In the next chapter, we employ our findings and a new model to devise an algorithm that fine-tunes the tradeoff towards reliable service for delay-sensitive applications.




Figure 5.11 – Short-term fairness and throughput obtained by simulation for parameters $CW_{min} = 8$, $m = 6$ and various values of d_0 and f . The deferral counter determines a tradeoff between throughput and fairness in 1901.




Figure 5.12 – Short-term fairness and throughput obtained by simulation with parameters $CW_{min} = 8$, $d_0 = 0$, $f = 2$ and various values of m .

	d_0	f	m	CW_{min}
small	T ↗	T ↗	T ↘	T ↗ if N is small
	F ↘	F ↘	F ↗	F →
large	T ↘	T ↘	T ↗	T ↗ if N is large
	F ↗	F ↗	F ↘	F →

Table 5.1 – Summary of the qualitative effects of each parameter on throughput (“T”) and short-term fairness (“F”).




Figure 5.13 – We study a testbed trace of $5 \cdot 10^4$ successfully transmitted packets for both 1901 and 802.11a, when two saturated stations A and B contend for the medium. The autocorrelation function of the identification of station that transmits shows that 1901 is less fair than 802.11.

5.7 Fairness and Coupling between the Stations

We explore the consequences of unfairness on modeling hypotheses that assume independence between the backoff processes of the stations. Bianchi [15] introduces such an approximation to study 802.11. His approximation works well for any number of stations. However, in Chapter 4, we observe that our similar assumption for 1901 might yield slight inaccuracies for a small number of stations. In the introductory example of this chapter and Figure 5.1, we learn the reasons of this inaccuracy.

The decoupling assumption relies on the approximation that the backoff processes of the stations are independent and that, as a consequence, stations experience the same time-invariant collision probability, independently of their own state and of the state of the other stations. In addition, to analyze 1901, it has been assumed that a station senses the medium busy with the same time-invariant probability (equal to the collision probability) at any time slot. We further show that the deferral counter introduces some coupling among the stations: After a station gains access to the medium, it can retain it for many consecutive transmissions before any other station can transmit. As a result, the collision and busy probabilities are not time-invariant for 1901 networks, which makes the decoupling assumption questionable.

In Figure 5.13, we present real evidence of the 1901 unfairness from a testbed trace of 50,000 frame transmissions. As we observe in Section 5.4, the worst unfairness for 1901 occurs for two stations. In Figure 5.13, we show the autocorrelation function of the identification of station that transmits in a network of two saturated stations A and B . Let X_i be the variable that indicates which station transmits successfully at the i th transmission. If A transmits, we take $X_i := 1$ and if B transmits, we have $X_i := 2$. We show the autocorrelation function of X_i , $1 \leq i \leq 5 \cdot 10^4$. Observe that it is positive for 1901 at lags smaller than 15, which means that if $X_i = 1$ for some i , it is likely that




Figure 5.14 – Simulations of 1901 (with CA1 parameters) and 802.11 for $N = 2, 5, 10$. Points show the collision probabilities at different backoff stages for all stations, and lines represent the solution of the fixed-point equations for the collision probability from decoupling-assumption models. The decoupling assumption is valid for 802.11, even for $N = 2$; whereas the collision probability depends on the backoff stage for 1901. More than 6000 samples have been obtained for each experiment.

$X_{i+1} = 1$. In contrast, for 802.11a we have a negative value of autocorrelation at lag 1 and a positive one at lag 2, which means that if $X_i = 1$ for some i , it is very likely that $X_{i+1} = 2$ and $X_{i+2} = 1$. The dependence of the non-transmitting station to the one monopolizing the medium questions the decoupling assumption, especially for such situations of severe unfairness.

For 802.11, the decoupling assumption has been shown to be valid for various settings as $N \rightarrow \infty$ [54]. In addition, it also works well for small numbers of stations [15, 55]. For 1901, the coupling induced by the deferral counter makes the collision probabilities state-dependent, which penalizes models based on the decoupling assumption when N is small. To see this, we plot on Figure 5.14 the collision probabilities experienced by 802.11 and 1901 stations, as a function of the backoff stage (i.e., as a function of the stations' state). On the same figure, we also show the collision probabilities computed with our model introduced in Chapter 4. Let O_k be the sequence of outcomes of attempted transmissions, i.e. $O_k \doteq 1$ if the k th transmission attempt results in a success, and $O_k \doteq 0$ when the outcome is a collision. The decoupling assumption asserts that the sequence $\{O_k\}$ consists of independent and identically distributed (i.i.d.) random variables. In Figure 5.14, we observe that for 1901, $\{O_k\}$ cannot be considered as i.i.d., because the collision probability observed at different backoff stages is not the same. Thus, the collision probability depends on the previous transmission attempts or on other stations' activity (both of which have an effect on the station's current backoff stage). In fact, the collision probability for 1901 increases with the backoff stage i , as shown in Figure 5.14. To address this issue, we introduce a model that does not rely on the "time-invariant collision-probability" assumption in the next chapter. Yet, note that, although the decoupling assumption might be a crude approximation for a small number of stations, it is expected to be valid for large N , due to similar reasoning as for 802.11. It has been proven analytically that the 802.11 decoupling assumption is valid as $N \rightarrow \infty$ [54, 56].

5.8 Related Work

As we are the first to delve into the short-term dynamics of 1901, there is not any related work on this MAC layer. Therefore, here we present related research on 802.11; we rely on 802.11 research to study 1901 in this chapter.

There are numerous studies on the fairness of MAC protocols in single collision-domain networks. First, Koksal et al. [52] explore the fairness of WaveLAN CSMA/CA, which works similarly to 1901, but without the need of the deferral counter: the WaveLAN stations *always* double their contention window after sensing the medium busy. The authors employ various metrics to study fairness, such as Jain’s fairness index with the sliding window method, and they conclude that this CSMA/CA is short-term unfair. Second, Berger-Sabbatel et al. [49] study the short-term fairness of 802.11 and introduce the number of inter-transmissions as a metric. They find that 802.11 does not exhibit short-term unfairness. Third, Bredel and Fidler [50] derive a closed-form expression for the conditional probability that a contending station transmits k packets, given that a tagged station transmits l packets within the same time interval; this study employs the distribution of inter-transmissions for studying fairness over long timescales. The authors support their analytical results by measurements.

We would like now to turn our attention to studies that focus on the coupling between the stations and on the relevant modeling-assumptions. In 2000, Bianchi [15] introduced one of the most popular models for 802.11 by relying on the decoupling assumption. Years after, many works [54, 56] validated Bianchi’s assumption for large number of stations. Surprisingly, this approximation works well even for two contending stations. Huang et al. [55] verify in practice the decoupling approximation for small number of stations by using a testbed of 802.11 stations and ns-2 simulations. Notably, they observe that the hypothesis is exact under greedy scenarios. They go beyond these scenarios and explore more modeling assumptions in non-saturated settings, 802.11e and mesh 802.11s networks. They find that common assumptions employed to incorporate station queues into the model are erroneous.

5.9 Summary

In this chapter, we study the fairness characteristics of 1901 and give a comparison with 802.11. 802.11 is conservative and employs large contention-window sizes to avoid collisions, whereas 1901 is more aggressive: It tries a small contention-window values first, but later engages a complex mechanism with the deferral counter to circumvent the high collision rate and to increase throughput. This extra complexity of 1901 comes at a price of fairness and delay. In terms of fairness, we identify two regimes: When the number of stations is $N \leq 15$, 802.11 is fairer than 1901. For $N \geq 15$, 1901 is fairer than 802.11. We show analytically that 1901 is significantly less fair than 802.11 when $N = 2$. Our findings imply that 1901 can be detrimental to delay-sensitive traffic, such as voice or video, especially when the default settings are used with priority CA1 of the IEEE

1901 standard. This suggests that end-user experience heavily relies on the existence of proper traffic-classification mechanisms that judiciously assign priorities to traffic flows.

Furthermore, we study the fairness implications on network performance and we better understand the effect of short-term unfairness on jitter. We find that the deferral counter in 1901 determines a tradeoff between throughput and fairness (hence, jitter), which is crucial for both best-effort and delay-sensitive traffic. *In 1901, any configuration providing throughput improvements results in deterioration of fairness.* We extensively explore the protocol parameters that affect this tradeoff. To the best of our knowledge, this is the first study on 1901 fairness and performance tradeoffs.

At a more general level, we extend the understanding of a commonly adopted modeling hypothesis for the MAC layer, as our research demonstrates that the decoupling assumption should be handled with care when analyzing MAC layers where there exists short-term unfairness. Although this assumption is proven analytically and experimentally to be viable for 802.11, we find that for 1901 it is stronger and does not hold when there are few stations contending for the medium. We attribute the inaccuracy of the hypothesis to the strong coupling between the stations, which is also related to short-term fairness.

6 Coupled Analysis of Multi-user Performance

6.1 Introduction

As we observed in the previous chapter, the introduction of the deferral counter in the PLC CSMA/CA induces a certain level of coupling between the stations, which penalizes the accuracy of models based on the *decoupling assumption*. This assumption was originally proposed in the 802.11 analysis of [15] and has been used in all works that have analyzed the 1901 CSMA/CA procedure so far (i.e., [36, 41] and our decoupling-assumption model presented in Chapter 4). In this chapter, we show that modeling accuracy can be substantially improved by avoiding this hypothesis.

We propose a theoretical framework for modeling the CSMA/CA process of 1901 without relying on the decoupling assumption. Recall that the decoupling assumption yields a time-invariant collision probability for any tagged station, which is independent of the network state. We now introduce a model, called *Drift*, that considers the coupling between stations and accurately captures 1901 performance: *The collision probability of the station depends on its own state and that of the other stations* (the state is typically the backoff stage). The model studies the evolution of the expected change (drift) in the number of stations at each backoff stage, and it is relatively compact: computing the throughput of the network only requires solving a system of m non-linear equations, where m is the number of backoff stages (the default value for 1901 is $m = 4$). We prove that this system of m equations admits a unique solution for a wide range of configurations. Compared to the decoupling analysis, this coupled model is more complex, as it includes more system states and it considers the states of all stations¹.

We investigate the accuracy of the Drift model and that of the decoupling-assumption model; Drift is the first model for 1901 reaching this level of high accuracy. In addition to accurate results in the stationary regime, Drift is the first model of 1901 that is able to analyze the transient regime, given the initial state of the system. This is useful for (i) evaluating convergence times to stationary regime, which depicts the ability of the protocol to adapt to dynamically changing conditions, (ii) investigating the global

¹The decoupling-assumption model studies only one tagged station given the independence assumptions.

asymptotic stability of the unique solution, that is, guaranteeing convergence to the same stationary point independently of the initial state of the system, and (iii) exploring short-term dynamics of the CSMA/CA process, such as performance metrics at each backoff stage. By exploiting our Drift analysis, we prove the global asymptotic stability of the unique solution for $m = 2$, and formulate the basis for proving stability for general values of m . We also leverage the high accuracy of the Drift model for to devise configurations that meet specific quality-of-service requirements.

The remainder of this chapter is organized as follows. In Section 6.2, we introduce our model, we prove that there is a unique solution for a wide range of configurations, and we discuss the global stability of the unique solution. Then, in Section 6.3, we present a thorough evaluation of the Drift model, comparing it with the decoupling-assumption one. In the same section, we investigate the transient regime of the model and the implications of configurations that yield more than one solution. In Section 6.4, we exploit the high-accuracy of the Drift model to design an algorithm that tunes the tradeoff between throughput and fairness, uncovered in the previous chapter. This tradeoff is particularly important for networks with small number of stations, where the Drift model is very accurate. We present related work on analyses without the decoupling assumption and on justifying its validity in Section 6.5. Finally, we summarize this chapter in Section 6.6.

6.2 Analysis

In this section, we introduce our coupled model for the 1901 CSMA/CA protocol (without relying on the decoupling assumption). Our analysis relies on the network assumptions introduced in Section 4.2.1, that is, perfect sensing and channel, N saturated stations in a single collision domain, infinite retry-limit, and homogeneous network.

6.2.1 The Drift Model

We model the network as a dynamical system that is described by the expected change in the number of stations at each backoff stage between any two consecutive time slots. In the stationary regime, the expected number of stations at each backoff stage is constant, hence we can compute performance by finding the equilibrium of the dynamical system.

Let us now introduce the variables of our model. Let m be the number of backoff stages and let $n_i, 0 \leq i \leq m - 1$ denote the number of stations at backoff stage i . Note that $\sum_{i=0}^{m-1} n_i = N$ and $n_i \in \mathbb{N}$. Let us further denote with τ_i the transmission probability at stage i , i.e., τ_i is the probability that a station at backoff stage i transmits at any given time slot. In addition, for a given station at backoff stage i , we denote with γ_i the probability that at least one other station transmits². We also denote with p_e the probability that no station transmits (or equivalently, that the medium is idle). Under the

²Compared to the decoupling-assumption model, where the collision probability γ is time invariant, here the collision probability depends on the station's state. To this end, in this chapter, we use the notation γ_i for the collision probability at backoff stage i .

assumption of independence of the transmission attempts, we have $p_e = \prod_{k=0}^{m-1} (1 - \tau_k)^{n_k}$, therefore

$$\gamma_i = 1 - \frac{p_e}{1 - \tau_i} = 1 - \frac{1}{1 - \tau_i} \prod_{k=0}^{m-1} (1 - \tau_k)^{n_k}. \quad (6.1)$$

Station Model: We now model the behavior of a given station at backoff stage i . We assume that the event that some other station transmits in a slot occurs with a constant probability γ_i , independent of the station's backoff and deferral counters values³. Hence, this is the probability that a transmission of the given station collides, as well as the probability that the station senses a slot busy when it does not transmit. The rest of the backoff process of the station is modeled accurately as a function of γ_i , drawing the station's backoff counter from a uniform distribution. With this model, we derive the probability that a station transmits and that it moves to the stage $i + 1$ due to the deferral counter at any random time slot. These two probabilities are used in the network model presented later.

By following similar reasoning as in Section 4.2.3, we can compute x_k^i , the probability that a station at backoff stage i jumps to the next stage $i + 1$ in k or fewer time slots due to sensing the medium busy, directly from γ_i . The only difference with the decoupling-assumption model introduced in Section 4.2.3 is that now, the collision probability depends on the station's state i . This yields

$$x_k^i = \sum_{j=d_i+1}^k \binom{k}{j} \gamma_i^j (1 - \gamma_i)^{k-j}. \quad (6.2)$$

Similarly to Section 4.2.3, we compute and use bc_i for the expected number of time slots spent by a station at backoff stage i , which now is a function of γ_i :

$$bc_i = \frac{1}{CW_i} \sum_{k=d_i+1}^{CW_i-1} \left[(k+1)(1 - x_k^i) + \sum_{j=d_i+1}^k j(x_j^i - x_{j-1}^i) \right] + \frac{(d_i+1)(d_i+2)}{2CW_i}. \quad (6.3)$$

Now, the transmission probability τ_i can be expressed as a function of x_k^i and bc_i , using the renewal-reward theorem, with the number of backoff slots spent in stage i being the renewal sequence and the number of transmission attempts (i.e., 0 or 1) being the reward. The expected number of transmission attempts at stage i can be computed similarly to bc_i . By dividing the expected number of transmission attempts at stage i with the expected time slots spent at stage i , τ_i is given by

$$\tau_i = \frac{\sum_{k=d_i+1}^{CW_i-1} \frac{1}{CW_i} (1 - x_k^i) + \frac{d_i+1}{CW_i}}{bc_i}. \quad (6.4)$$

³With this assumption, we are neglecting the coupling between the deferral counter decrements of different stations. Note, however, that this does not couple the actual transmissions, as these follow a separate random process; as a result, the coupling due to the deferral counter is somehow diluted.

Notation	Definition (at backoff stage i , $0 \leq i \leq m - 1$)
n_i	Number of stations
γ_i	Probability that at least one other station transmits at any slot
p_e	Probability that the medium is idle at any slot (independent of i)
x_k^i	Probability that a station leaves stage i due to sensing the medium busy $d_i + 1$ times during k slots
bc_i	Expected number of backoff slots
τ_i	Probability that a station transmits at any slot
β_i	Probability that, at any slot, a station leaves stage i due to sensing the medium busy $d_i + 1$ times
F_i	Expected change in n_i between two consecutive slots
\bar{n}_i	Expected number of stations

Table 6.1 – Notation list relevant to a station at backoff stage i .

Similarly, we define β_i as the probability that, at any given slot, a station at stage i moves to the next backoff stage because it has sensed the medium busy $d_i + 1$ times. It can be easily seen that β_i is given by

$$\beta_i = \frac{\sum_{k=d_i+1}^{CW_i-1} \frac{1}{CW_i} \sum_{j=d_i+1}^k (x_j^i - x_{j-1}^i)}{bc_i}. \quad (6.5)$$

It will be very important in the following to remember that τ_i and β_i are functions of γ_i (through x_k^i and bc_i). Our notation is summarized in Table 6.1. We next study the evolution of the expected change (drift) in the number of stations at each backoff stage i .

Transient Analysis of the System

Building on the analysis above, we now introduce the Drift model. A key feature is that we do not assume that the stations are decoupled, as the collision probability is allowed to depend on the station's state. To study the system, we use a vector that includes the number of stations at each backoff stage. In particular, let $\mathbf{X}(t) = (X_0(t), X_1(t), \dots, X_{m-1}(t))$ represent the number of stations at each backoff stage $(0, 1, \dots, m - 1)$ at time slot t . We use the notation $\mathbf{n}(t) = (n_0(t), n_1(t), \dots, n_{m-1}(t))$ to denote a realization of $\mathbf{X}(t)$ at some time slot t .

Network Model: To model the network, we rely on the simplifying assumption that a station transmits, or moves to the next backoff stage upon expiring the deferral counter, with a constant probability (independently of previous time slots). This is necessary as otherwise, we would need to keep track of the backoff and deferral counter values of each station and the model would become intractable. In particular, our assumptions are as follows: (i) a station at backoff stage i attempts a transmission in each time slot with a constant probability $\tau_i(\gamma_i)$; and (ii) a station at backoff stage i moves to backoff stage $i + 1$ due to the deferral counter expiration with a constant probability $\beta_i(\gamma_i)$ in each time slot where it does not transmit. Both τ_i and β_i depend on the probability γ_i that

the station senses a slot busy, which is computed from the transmission probabilities of the other stations following (6.1).

With the above assumptions, $\mathbf{X}(t)$ is a Markov chain. The transition probabilities τ_i and β_i depend on the state vector $\mathbf{n}(t)$ and they can be computed from (6.1), (6.4) and (6.5); hereafter, to simplify notation, we drop the input variable t from $\gamma_i(t)$, $\tau_i(t)$, $\beta_i(t)$, and $\mathbf{n}(t)$ as the equations are expressed for any slot t .

Let now $\mathbf{F}(\mathbf{n}) = \mathbb{E}[\mathbf{X}(t+1) - \mathbf{X}(t) | \mathbf{X}(t) = \mathbf{n}]$ be the expected change in $\mathbf{X}(t)$ over one time slot, given that the system is at state \mathbf{n} . Function $\mathbf{F}(\cdot)$ is called the *drift* of the system, and is given by

$$F_i(\mathbf{n}) = \begin{cases} \sum_{k=1}^{m-1} n_k \tau_k (1 - \gamma_k) - n_0 \tau_0 \gamma_0 - n_0 \beta_0, & i = 0 \\ n_{i-1} (\tau_{i-1} \gamma_{i-1} + \beta_{i-1}) - n_i (\tau_i + \beta_i), & 0 < i < m - 1 \\ n_{m-2} (\tau_{m-2} \gamma_{m-2} + \beta_{m-2}) - n_{m-1} \tau_{m-1} (1 - \gamma_{m-1}), & i = m - 1. \end{cases} \quad (\text{DRIFT})$$

(DRIFT) is obtained by balancing, for every backoff stage, the average number of stations that enter and leave this backoff stage. In particular, n_0 increases by 1 only when some station transmits successfully. Since such a station could be in any of the other backoff stages and there are n_k stations in stage k , this occurs with probability $\sum_{k=1}^{m-1} n_k \tau_k (1 - \gamma_k)$. Similarly, n_0 decreases when some stations at stage 0 are either involved in a collision (which occurs with probability $n_0 \tau_0 \gamma_0$), or do not transmit and sense the medium busy $d_0 + 1$ times (which occurs with probability $n_0 \beta_0$). The decrease of n_0 in both cases is 1, thus the expected decrease is equal to the sum of the two probabilities. The resulting drift F_0 is computed by adding all these (positive and negative) expected changes in n_0 .

Similarly, F_i , $0 < i < m - 1$ is computed by observing that in these backoff stages, n_i decreases if and only if some stations at stage i sense the medium busy or transmit. n_i increases if and only if some stations at stage $i - 1$ sense the medium busy or transmit and collide. Finally, n_{m-1} increases after some stations at stage $m - 2$ experience a collision or sense the medium busy $d_{m-2} + 1$ times. It decreases only after a successful transmission at stage $m - 1$.

The evolution of the expected number of stations $\bar{\mathbf{n}}(t) \doteq \mathbb{E}[\mathbf{X}(t)]$ is described by the m -dimensional dynamical system

$$\bar{\mathbf{n}}(t+1) = \bar{\mathbf{n}}(t) + \mathbf{F}(\bar{\mathbf{n}}(t)), \quad (\text{DYNSYS})$$

where $\mathbf{F}(\bar{\mathbf{n}}(t))$ is given by (DRIFT).

Our model relies on the key insight that the stochastic system stays close to the typical state given by the equilibrium of (DYNSYS), and accurate estimates of various metrics such as throughput can be obtained by assuming that the system is in this typical state at all times. However, in reality the stochastic system might stay in other states with a certain probability. It is intuitive that the model becomes more accurate as the

number of stations in the system grows: If the number of stations at a each backoff stage is very large, it is expected that the behavior of the stochastic system is close to the deterministic one given by (DYNSYS) due to the law of large numbers. This has been proven by Sharma et al. [57] for 802.11: By analyzing a properly scaled version of the stochastic system (where scaling is in time as well as in magnitude), the authors prove a functional “law of large numbers” for the stochastic system, showing that the stochastic process converges to the deterministic one as $N \rightarrow \infty$. The main insights behind the proof of Sharma et al. [57] go beyond the specifics of the 802.11 protocol, hence the same result applies to 1901. We prove this key insight for 1901 in Section 6.2.3 and further discuss these convergence results in Section 6.5. Although this key insight is proven for large N , in Section 6.3, we observe that the Drift model is very accurate for N as small as two.

Steady-State Analysis of the System

In order to know the average number of stations at each backoff stage at steady state, we can compute the equilibrium point(s) of system (DYNSYS), which is the stationary regime where the average number of stations at each backoff stage remains constant and which is studied in this subsection. This information enables us to compute actual throughput figures. To compute the equilibrium point(s) of (DYNSYS), we impose the condition $\mathbf{F}(\bar{\mathbf{n}}) = \mathbf{0}$, which yields

$$\bar{n}_i = \left(\frac{\tau_{i-1}\gamma_{i-1} + \beta_{i-1}}{\tau_i + \beta_i} \right) \bar{n}_{i-1}, \quad 1 \leq i \leq m-2, \quad \bar{n}_{m-1} = \left(\frac{\tau_{m-2}\gamma_{m-2} + \beta_{m-2}}{\tau_{m-1}(1 - \gamma_{m-1})} \right) \bar{n}_{m-2}.$$

Let us define

$$K_0 \doteq 1, \quad K_i \doteq \frac{\tau_{i-1}\gamma_{i-1} + \beta_{i-1}}{\tau_i + \beta_i}, \quad 1 \leq i \leq m-2, \quad K_{m-1} \doteq \frac{\tau_{m-2}\gamma_{m-2} + \beta_{m-2}}{\tau_{m-1}(1 - \gamma_{m-1})}. \quad (6.6)$$

Since $\sum_{i=0}^{m-1} \bar{n}_i = N$, the equilibrium $\hat{\mathbf{n}}$ of (DYNSYS) is given by

$$\hat{n}_i = \frac{N \prod_{j=0}^i K_j}{\sum_{k=0}^{m-1} \prod_{j=0}^k K_j}, \quad 0 \leq i \leq m-1. \quad (\text{EQ})$$

Recall that τ_i and β_i are functions of γ_i , given by (6.4) and (6.5). Thus, the \hat{n}_i 's in (EQ) are also functions of γ_i , $0 \leq i \leq m-1$. From the above, substituting (EQ) in (6.1) yields a system of m equations with m unknowns γ_i for $0 \leq i \leq m-1$.

Throughput Evaluation

After solving the equations for finding the steady-state number of nodes $\hat{n}_0, \dots, \hat{n}_{m-1}$ at each backoff stage, we can compute the throughput of the network as follows. The probability that a slot is idle is p_e . The probability of a successful transmission of a station at stage i is $\tau_i(1 - \gamma_i)$. Therefore, the probability p_s that a slot contains a

successful transmission is given by $p_s = \sum_{i=0}^{m-1} \hat{n}_i \tau_i (1 - \gamma_i)$, assuming that \mathbf{n} remains in a neighborhood of the equilibrium point $\hat{\mathbf{n}}$. Let p_c denote the probability that a slot contains a collision. We have $p_c = 1 - p_e - p_s$. We now have enough information to compute the normalized throughput S of the network as

$$S = \frac{p_s D}{p_s T_s + p_c T_c + p_e \sigma}, \quad (6.7)$$

where D is the frame duration, T_s is the duration of a successful transmission, T_c is the duration of a collision, and σ is the time slot duration. In Section 6.3, we show that the Drift model is very accurate for a wide range of configurations.

6.2.2 Uniqueness of the Equilibrium Point

In this subsection we prove that, as long as the configuration of CW_i 's and d_i 's is chosen such that the sequence τ_i is decreasing with i for any n_i distribution, then the equilibrium point given by (EQ) is unique. We argue that such a condition should be met by any sensible configuration of CW_i 's and d_i 's. The argument is as follows. Jumping to the next backoff stage is an indication of high contention, either because of a collision or a sequence of busy slots. Therefore, in this case τ_i should decrease with i , and the high contention should be dissolved by reducing the aggressiveness of the sources. Note that similar studies for the 802.11 MAC protocol [30, 54] require the same sufficient condition (i.e., τ_i decreasing with i) for the model to admit a unique solution. To simplify the exposition, we define this condition as follows.

$$\tau_i > \tau_{i+1}, 0 \leq i \leq m - 2. \quad (\text{COND})$$

Theorem 2 states that that if (COND) is satisfied, the equilibrium point given by (EQ) is unique. The proof is given in Appendix A.2.

Theorem 2. *If (COND) is satisfied, then the system of equations formed by (EQ) and (6.1) for $0 \leq i \leq m - 1$ has a unique solution.*

We next provide some configuration guidelines for (CW_i, d_i) that ensure that (COND) is satisfied. In Section 6.3.2, we discuss a counterexample of a configuration that does not satisfy (COND) and does not yield a unique solution to the system of equations formed by (EQ) and (6.1).

Protocol Configurations Satisfying (COND)

Before showing in Theorem 3 that (COND) is satisfied for a wide range of configurations, we introduce a useful lemma. Note that compared to 802.11, where τ_i is a function of only CW_i , the analysis here is substantially more challenging, because τ_i is a function of CW_i , d_i , and γ_i .

We have to investigate the relationship between τ_i and τ_{i+1} . Recall that these two

transmission probabilities are functions of two different collision probabilities γ_i and γ_{i+1} , respectively, which makes the analysis challenging. Assume that the collision probability is the same for two successive backoff stages i and $i + 1$, and is equal to γ_i . Under this hypothesis, in Lemma 3, we show that if $\tau_i(\gamma_i) > \tau_{i+1}(\gamma_i), \forall \gamma_i \in [0, 1]$, then $\tau_i(\gamma_i) > \tau_{i+1}(\gamma_{i+1})$, for any pair γ_i, γ_{i+1} that satisfies (6.1). In Theorem 3, we provide some sufficient conditions to guarantee that $\tau_i(\gamma_i) > \tau_{i+1}(\gamma_i)$ is satisfied for all $\gamma_i \in [0, 1]$ and $0 \leq i < m - 1$; from Lemma 3, this implies that (COND) is satisfied. The proofs are provided in Appendix A.2.

Lemma 3. *Let γ_i^s be a value of the collision probability at stage i . Then, if $\tau_i(\gamma_i^s) > \tau_{i+1}(\gamma_i^s)$ for all $\gamma_i^s \in [0, 1]$, we have $\tau_i(\gamma_i) > \tau_{i+1}(\gamma_{i+1})$ for any n_i distribution.*

The following theorem provides some sufficient conditions on the (CW_i, d_i) configurations that ensure that (COND) holds. Note that Lemma 3 could be employed to show that (COND) holds for more configurations than the ones covered by the theorem; indeed, it is sufficient to show that the configuration satisfies the hypothesis of the Lemma 3 for all $0 \leq i \leq m - 2$.

Theorem 3. (COND) holds if the following condition is satisfied for $0 \leq i \leq m - 2$

$$CW_{i+1} > \begin{cases} CW_i, & \text{if } d_{i+1} = d_i \\ 2CW_i - d_i - 1, & \text{otherwise.} \end{cases} \quad (6.8)$$

Observe that, from Table 2.2, the above constraints (6.8) on CW_i and d_i are compliant with the IEEE 1901 standard, except for the class CA2/CA3 at backoff stage $i = 1$. The results obtained in this chapter suggest that it might be worth revisiting the configuration of this priority class; indeed, for the proposed configuration of CA2/CA3 we have $\tau_2 > \tau_1$ and (COND) does not hold. We next discuss whether (COND) is sufficient for the global asymptotic stability of (EQ).

6.2.3 Convergence to the Unique Equilibrium Point

In addition to showing that the system (DYNSYS) has only one equilibrium point, it is also interesting to show that it converges to this equilibrium point, from all possible initial states, i.e., that (EQ) is globally asymptotically stable. Such a proof is important for the rigorous explanation of the accuracy of the Drift model in the stationary regime, and the corroboration of the decoupling assumption introduced in Chapter 4, as N becomes large. In fact, if the equilibrium point is unique but not stable, and there exists a stable periodic solution in the stationary regime, then the decoupling assumption does not hold: The collision probability actually oscillates due to the periodic solution and it cannot be modeled as time invariant [58].

Notably, the authors of the Drift model for 802.11 [57] rely on an asymptotic system as $N \rightarrow \infty$ to prove global asymptotic stability for $m = 2$. To this end, they analyze a

scaled version of $\mathbf{X}(t) = (X_0(t), X_1(t), \dots, X_{m-1}(t))$, that is, $Y_i^{(N)}(t) = X_i(\lfloor Nt \rfloor)/N$, for all $0 \leq i \leq m-1$, where scaling is performed both in time and in magnitude. They also scale the transition probabilities per station (i.e., a station at backoff stage i transmits with probability τ_i/N), to prevent the collision probability to converging to 1, as N gets large. They first prove that $Y_i^{(N)}(t)$ converges to a deterministic process given by the solution of system of non-linear differential equations given by the scaled version of (DYNSYS). This validates that, as $N \rightarrow \infty$, the random variable $\mathbf{X}(t)$ stays close to a deterministic process, which is the key insight of our Drift model introduced in Section 6.2.1. Then, the authors prove that the system of non-linear differential equations has a unique equilibrium point (independently of N) and that this equilibrium is globally asymptotically stable for $m = 2$ (the proof for $m > 2$ is challenging even for 802.11).

Let us now discuss the specifics of the 1901 Drift model with respect to asymptotic convergence and global asymptotic stability. As in all the previous analyses of 802.11 [54, 56, 57], we consider a scaled version of our system, $\mathbf{X}(\lfloor Nt \rfloor)/N$, where time is accelerated by a factor of N while the transition probabilities are scaled down by the same factor:

- By scaling time, the evolution of time slots is accelerated by N , such that a variable at time t before this operation is translated into the scaled one at time t/N .
- By scaling the transition probabilities, the evolution of each node is slowed down by a factor of N .

With this scaling, the expected change of the state of the system over two consecutive time-slots is order of $1/N$, which tends to zero as $N \rightarrow \infty$. By accelerating the evolution of time slots by N , the change of the system over time remains in the same order of magnitude as the original system.

Following the above, to scale down the transition probabilities we let the probability that a station attempts a transmission at backoff stage i be $a_i^{(N)}(\gamma_i) \doteq \tau_i(\gamma_i)/N$, and the probability that it jumps to the next backoff stage due to the expiration of the deferral counter be $b_i^{(N)}(\gamma_i) \doteq \beta_i(\gamma_i)/N$. We further define $y_i(t) \doteq n_i(t)/N$ as the fraction of stations at backoff stage i . With this, when $N \rightarrow \infty$ we obtain the following deterministic (asymptotic) system:

$$\frac{dy_i}{dt} = \phi_i = \begin{cases} \sum_{k=1}^{m-1} y_k \tau_k(\rho)(1 - \rho) - y_0(\tau_0(\rho)\rho + \beta_0(\rho)), & i = 0 \\ y_{i-1}(\tau_{i-1}(\rho)\rho + \beta_{i-1}(\rho)) - y_i(\tau_i(\rho) + \beta_i(\rho)), & 0 < i < m-1 \\ y_{m-2}(\tau_{m-2}(\rho)\rho + \beta_{i-1}(\rho)) - y_{m-1}\tau_{m-1}(\rho)(1 - \rho), & i = m-1. \end{cases} \quad (\text{ODE})$$

where $\rho = 1 - e^{-\sum_{k=0}^{m-1} y_k \tau_k(\rho)}$ is the probability that a given time slot is busy in the asymptotic system⁴, and $\phi_i = \lim_{N \rightarrow \infty} F_i(N \cdot \mathbf{y})$ is computed by (DRIFT) (after replacing the transition probabilities by $a_i^{(N)}(\gamma_i)$ and $b_i^{(N)}(\gamma_i)$).

The following theorem shows that the stochastic system under study converges

⁴The limit $\rho = \lim_{N \rightarrow \infty} \gamma_i$ is computed given the limit $\lim_{N \rightarrow \infty} (1 - x/N)^N = e^{-x}$.

to the deterministic model given above as $N \rightarrow \infty$, which confirms that the proposed analysis becomes very accurate as the number of stations grows large. The proof is based on validating simple conditions on then smoothness of the transition probabilities and the drift (DRIFT) and on the assumption that “the coefficient of variation of the number of objects that do a transition in one time slot remains bounded for large N ”, which are rigorously established in [58]. We give the proof in Appendix A.2.

Theorem 4. *As $N \rightarrow \infty$, the 1901 scaled random system $\mathbf{X}(\lfloor Nt \rfloor)/N$ converges to the deterministic process $\mathbf{y}(t)$ given by (ODE).*

In the following theorem, we show that the system (ODE) has a globally asymptotically stable equilibrium point for $m = 2$. The proof can be found in Appendix A.2.

Theorem 5. *If (COND) is satisfied, the system (ODE) is globally asymptotically stable for $m = 2$.*

Despite the intuition from previous work on 802.11 [54, 57], we cannot yet prove global stability of the complex 1901 ODE for general values of m . Yet, based on our extensive simulations, we conjecture that (COND) yields global asymptotic stability of (EQ). To corroborate the above conjecture and to confirm the global asymptotic stability of the dynamical system given by (DYNSYS), we conduct an additional comprehensive numerical study. Specifically, we evaluate (DYNSYS) for a wide range of configurations satisfying (COND), comprising all the following values for the various parameters: $CW_i = \{8, 16, 32, 64\}$, $m = \{3, 4, 5, 6\}$ and $d_i = \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 20, 25, 30\}$. For each configuration, we choose randomly 100 different initial points $\mathbf{n}(0)$ and evaluate the trajectory of (DYNSYS) until it converges (with an error of 10^{-8}). In total, around 10^6 tests of convergence are conducted and, in all tests, (DYNSYS) converges to the equilibrium given by (EQ). Based on the evidence provided by this numerical study, the intuitive arguments exposed above, and by the theoretical results for some specific cases (such as 802.11, $m = 2$ and $N \rightarrow \infty$), we conjecture that (DYNSYS) converges to the unique equilibrium point for any configuration satisfying (COND).

In Section 6.5, we review related work on the global asymptotic stability of similar 802.11 models and we discuss the necessary steps for similar proofs of 1901. It turns out, that such 802.11 models have been used to justify the decoupling approximation originally introduced by Bianchi [15] and to verify that there are no oscillations (periodic solutions) in the stationary regime. The decoupling hypothesis has been rigorously validated under various 802.11 configurations as N becomes large. In all cases, the proof relies on a scaled system that transforms to a continuous-time dynamical system as $N \rightarrow \infty$.

6.3 Performance Evaluation

In this section, we evaluate the performance of 1901 under different configurations and scenarios by using simulations. Our simulator is validated experimentally in Section 4.4.1. We consider the timing parameters and assumptions presented in Section 4.4.2.

We compare the accuracy of the Drift model against the decoupling-assumption (D.A.) model introduced in Chapter 4 and we summarize the models differences. Furthermore, we evaluate other aspects of the Drift model, such as the performance of the configurations that do not satisfy (COND) and the accuracy of the model in transient regime. We conduct simulations for a wide range of configurations (CW_i, d_i, m).

6.3.1 Drift vs. Decoupling-Assumption Model

We now compare the Drift model with the D.A. model for various configurations and number of stations. In Figure 6.1, we show the throughput obtained by 1901 with the default parameters for the two priority classes CA1 and CA3 (CA0 and CA2 are equivalent). We also show the throughput predicted by the two models. The model based on the decoupling assumption is substantially less accurate for CA1 when N is small, because the class CA1 uses larger contention windows, which increases the time spent in backoff and, as a result, the coupling between stations.




Figure 6.1 – Throughput obtained by simulation, and with the Drift and D.A. models for the default configurations of 1901 given in Table 2.2.

We study the accuracy of the two models in more general settings. To this end, we introduce a factor f , such that at each stage i , the value of d_i is given by $d_i = f^i(d_0 + 1) - 1$. This enables us to define various sequences of values for the d_i 's, using only f and d_0 . At each stage i , CW_i is given by $CW_i = 2^i CW_{min}$, and there are m backoff stages ($i \in \{0, m - 1\}$). In Figure 6.2, we show the throughput for various such values of d_0 and f , with $CW_{min} = 8$ and $m = 5$. We observe that the D.A. model achieves good accuracy when the d_i 's are large, because in these configurations, the deferral counter is less likely to expire, which reduces the coupling among stations. Note that the Drift model achieves good accuracy when the d_i 's are small, while there is a small deviation for large d_i 's; this is due to the assumptions of our network model of Section 6.2.1, which are not used by the D.A. model⁵.

⁵The 802.11 model that does not rely on the decoupling assumption [57] has a similar deviation compared to Bianchi's model [15].




Figure 6.2 – Throughput obtained by simulation, and with the Drift and D.A. models for different configurations. The initial values d_i of the deferral counter at each backoff stage are given by $d_i = f^i(d_0 + 1) - 1$.

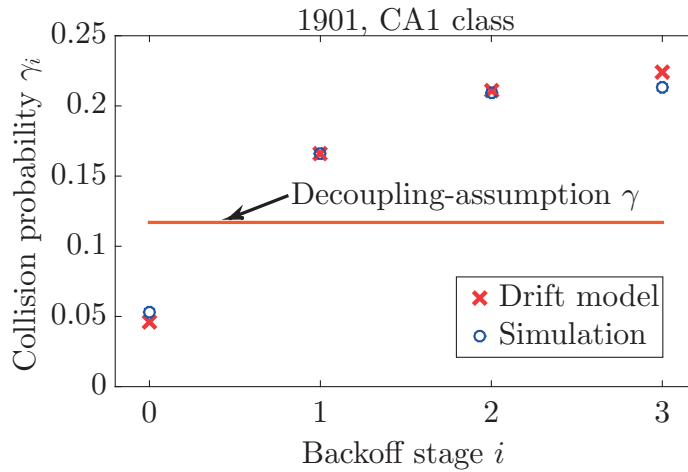


Figure 6.3 – Simulations and model evaluations of the collision probability with 1901 (with CA1 parameters) for $N = 2$. Drift model accurately predicts the collision probability at each backoff stage i , γ_i . For reference, we also show the collision probability γ predicted by the decoupling-assumption model.

In Figure 6.3, we evaluate the ability of the Drift model to predict the collision probability at each backoff stage i . We verify that indeed, the collision probability is not time invariant for the default class of PLC, hence also for configurations that introduce strong coupling between the stations. As we also observe in Section 5.7, this coupling penalizes the accuracy of the decoupling-assumption model. The Drift model accurately estimates the collision probabilities γ_i , $0 \leq i \leq 3$.

Drift Model	Decoupling-Assumption Model
Coupled model of all stations	Model of a tagged station given independence
System of m non-linear equations (complex)	Non-linear fixed-point equation (simple)
Both transient and stationary regimes	Only stationary regime
Enables short-term dynamics analysis	Only long-term average analysis
Enables future global-stability proofs	Nothing can be inferred about global stability
More accurate for small N , except when d_i 's are large	Much less accurate for small N , except when d_i 's are large
Application: accurate throughput prediction for tuning the tradeoff between throughput and fairness	Application: throughput enhancements for large N , thanks to simplicity

Table 6.2 – Comparison of the Drift and decoupling-assumption models. In addition to the differences described here, let us note that the models have the same accuracy for large N (typically $N \geq 15$).

We now summarize the differences between our two CSMA/CA performance models presented in this thesis. Table 6.2 compares in detail the two models. As we observe, both models have practical benefits and disadvantages, complementing each other. They can be used for different practical applications and modeling insights. In the rest of this section and in the next one, we discuss insights and applications of the Drift model.

6.3.2 Uniqueness of the Solution and Selected Counterexample

One of the fundamental results of our steady-state analysis is that there is a unique equilibrium for configurations that satisfy (COND). In this subsection we investigate the performance of the system depending on the (non-)unicity of the equilibrium of the dynamical system. To this end, we explore a counterexample of a configuration that does not satisfy (COND) and does not yield a unique equilibrium for the dynamical system (DYNSYS). An example of such a configuration, which yields 3 equilibrium points for $N = 10$, is the following:

$$\{CW_i, d_i\} = \begin{cases} \{32, 3\}, & 0 \leq i \leq 3 \\ \{4, \infty\}, & 4 \leq i \leq 53 \\ \{64, 3\}, & 54 \leq i \leq 59. \end{cases} \quad (6.9)$$

To study this configuration, we compute the instantaneous p_e , i.e, the probability that a time-slot is idle, for every 500 slots in simulation. Figure 6.4 shows the results for the CA1 class and for the configuration given by (6.9). We observe that for CA1 class, for which we have a unique equilibrium, the instantaneous p_e is approximately equal to the one given by the equilibrium point of the dynamical system (DYNSYS). However, for configuration (6.9), p_e oscillates between two of the equilibrium points and the value




Figure 6.4 – Simulation of a system with a unique equilibrium point (left, configuration of CA1 class) and with 3 equilibrium points (right, configuration given by (6.9)). The probability p_e is computed for each 500 slots and is shown for one simulation run (plain blue). The values of the equilibrium point(s) are also shown for each system (dashed red). The same behavior was observed for both systems for multiple, repeated simulation runs, not shown here.

of p_e averaged the entire simulation run is not equal to any of the equilibrium points; indeed the average p_e obtained by one simulation run is 0.3478, whereas the values of the equilibrium points of (DYNSYS) are $(p_e^1, p_e^2, p_e^3) = (0.5202, 0.2087, 0.0585)$ ⁶.

Our results show that the equilibrium points (EQ) is not sufficient to characterize the performance of the real system when (EQ) is not unique: The real system might oscillate and, as a result, the behavior might not be close to any of the equilibrium points. They also suggest that such configurations should be avoided as they might lead to an unstable thus, undesirable behavior. Indeed, multiple equilibria yield metastable regimes and typically involve severe unfairness or network collapse. For instance, with configuration (6.9) some stations remain at a backoff stage with $CW_i = 4$ for long periods, leading to a very high collision probability and low throughput. The performance problems resulting from metastable regimes are reported in [16, 54] for 802.11.

Configuration Guidelines with Respect to (COND)

As discussed in Section 6.2.2, (COND) is not only a condition for uniqueness, but also a configuration guideline for proper reaction to high contention. Jumping to the next backoff stage is an indication of high contention hence, to dissolve the current contention, the transmission aggressiveness should decrease, that is $\tau_{i+1} < \tau_i$. We now show that configurations where τ_i is increasing with i perform poorly. To confirm this, in the following we run several experiments with different values of f , where f now can be

⁶To search and find configurations with more than one equilibrium point, we converted the system (EQ) into a fixed-point equation with respect to p_e . The procedure is given in the proof of Theorem 2 in Appendix A.2. We find graphically the number of fixed-points and their approximate values. Then to obtain their exact value, we solve the system (EQ) with initial values close to the approximate solutions.




Figure 6.5 – Performance of 1901 with parameters $CW_i = \lfloor f^{i-1} \cdot 8 \rfloor$, and $d_i = \lceil f^{i-1} - 1 \rceil$ for different values of f . Lines represent throughput obtained by simulation and points show throughput computed by our model (left). We also present the ratio τ_1/τ_0 computed using our model (right).

a positive real number, and with $CW_{min} = 8$, $d_0 = 0$. Figure 6.5 presents throughput obtained by simulation and with our model for various values of f , with $CW_i = \lfloor f^{i-1} \cdot 8 \rfloor$, $d_i = \lceil f^{i-1} - 1 \rceil$, $0 \leq i \leq 3$. Results show that configurations with τ_i increasing, i.e., $f < 1$, yield poor performance. This supports our argument that (COND) should be met to ensure good performance. Theorem 3 provides some configuration guidelines to ensure that (COND) is satisfied.

As it can be seen from the figure, throughput performance improves for large f . However, a closer look at the protocol behavior for different f 's reveals that, while large f 's provide very good throughput performance, they also suffer from severe unfairness (e.g., $f = 16$). Indeed, for such configurations only one station grasps the channel, while the others move to higher backoff stages with much larger values of CW_i and barely transmit. This shows that throughput considerations are not sufficient to properly evaluate the suitability of a given 1901 configuration, and short-term fairness also needs to be taken into account. This is further explored in Section 6.4.

6.3.3 Accuracy of the Deterministic System in Transient Regime

The above experiments have focused on the accuracy of our steady-state analysis for the stationary regime. In the following, we investigate the accuracy of the analysis for the transient regime. To this end, we consider a system with $N = 20$ stations and two different configurations, and compare the expected number of stations $\bar{\mathbf{n}}(t)$ obtained from (DYNSSYS) and from simulations, as a function of the time slot t , when the initial condition at time slot 0 is $\mathbf{n}(0) = \{20, 0, 0, \dots, 0, 0\}$.

We focus on a configuration that follows the 1901 standard, i.e., CA1 class, and on a configuration that follows the 802.11 standard, i.e., the deferral counter does not expire. The results from the experiments described above are shown in Figure 6.6. We




Figure 6.6 – Convergence to the equilibrium point of the number of stations for backoff stages 1 and 3, for the CA1 class (left) and for a configuration with $d_i \rightarrow \infty$ and $CW_i = 2^i CW_0$, $\forall i$, $CW_0 = 8$, $m = 7$, (right). Both the expected values obtained from (DYNSYS) and the average values obtained from 2000 simulation runs are shown.

observe that our model works well both in terms of accuracy and of convergence times. As far as accuracy is concerned, there is slightly higher inaccuracy in the transient regime than in the stationary regime, which is due to the assumption on the constant transition probabilities β_i and τ_i ⁷. The convergence time to the equilibrium points is also captured by our model with reasonable accuracy. This time is higher for the 802.11 system for two reasons: (i) in the 802.11 system, the stations are allowed to have larger backoff counters and to move into higher backoff stages; (ii) in the 1901 system, the stations change their backoff stage with a higher probability than in 802.11 due to the deferral counter.

6.4 Adjusting the Fairness-Throughput Tradeoff with the Drift Model

The Drift model is the first one reaching this level of accuracy, especially for small number of stations which currently is a frequent scenario in practice. As we discussed in Chapter 5, under this scenario, there might exist strong unfairness in the network due to the deferral counter. The unfairness has an impact on delay-variance (jitter), which is crucial for delay-sensitive applications. In fact, there is a tradeoff between throughput and fairness that can be adjusted based on the protocol parameters d_i and m (see Table 5.1). Recall that the larger the d_i 's are (respectively, the lower the m), the fairer but less efficient the protocol is. So far, the lack of an accurate analysis for small

⁷To confirm that the deviations are due to this assumption, we simulated the Markov chain $\mathbf{X}(t)$ with constant transition probabilities, and verified that the trajectory of $\mathbf{X}(t)$ averaged over 300 runs coincides with the solution of (DYNSYS).

6.4. Adjusting the Fairness-Throughput Tradeoff with the Drift Model

number of stations prevented us from fully exploiting this tradeoff.

To exploit this tradeoff and to find efficient configurations that meet specific quality-of-service criteria, we can now use the findings summarized in Table 5.1, together with our Drift model. As an example, we propose a simple heuristic algorithm that finds an efficient configuration in terms of jitter, given an arbitrary throughput requirement (if such a configuration exists). Our method is detailed in Algorithm 1, and works as follows. It orders (by increasing order of values) the sets of possible values taken by d_0 and m in two sequences named \mathcal{D} and \mathcal{M} , respectively. The possible values of CW_{min} are stored in a vector named \mathcal{C} . It then performs a binary search on \mathcal{D} : for a given d_0 in \mathcal{D} , it tests all combinations of parameters (m, CW_{min}) (by increasing order of m). When such a configuration satisfies the throughput requirement, the algorithm stores it and tries a larger value for d_0 (as a larger d_0 can potentially yield better jitter). Conversely, if no configuration meeting the throughput requirement is found, the algorithm considers smaller values for d_0 (which yield higher throughput, potentially at the expense of jitter). The algorithm ends when it has found the best configuration with the largest possible d_0 , still satisfying the constraint.

Algorithm 1: 1901 configuration for minimum jitter

Input: Throughput requirement S , number of stations N , sequences \mathcal{D} , \mathcal{M} and \mathcal{C} of possible values for d_0 , m and CW_{min} , respectively

Output: A configuration (d_0, m, CW_{min}) that minimizes the jitter and provides throughput at least S , if it exists (returns `null` otherwise)

Initialize:

Sort the sequences \mathcal{M} and \mathcal{D} by increasing order of values

Set $h_1 \leftarrow 0$ and $h_2 \leftarrow |\mathcal{D}| - 1$

Set $config \leftarrow \text{null}$

while $h_1 \leq h_2$ **do**

 Set $break_flag \leftarrow \text{false}$

$h_3 \leftarrow \lceil (h_2 - h_1)/2 \rceil + h_1$

 Set $d_0 \leftarrow \mathcal{D}_{h_3}$ (i.e., the h_3 -th element of sequence \mathcal{D})

for each $m \in \mathcal{M}$ and $CW_{min} \in \mathcal{C}$ **do**

 evaluate throughput \hat{S} from model when using configuration (d_0, m, CW_{min})

if $\hat{S} \geq S$ **then**

 Set $config \leftarrow (d_0, m, CW_{min})$

 Set $h_1 \leftarrow h_3 + 1$

 Set $break_flag \leftarrow \text{true}$

break out of for loop

end

end

if $break_flag == \text{false}$ **then**

 Set $h_2 \leftarrow h_3 - 1$

end

end

Return: $config$

Because it employs a binary search, the complexity of this algorithm is $O(|\mathcal{C}| \cdot |\mathcal{M}| \cdot \log(|\mathcal{D}|))$. We evaluate it on the sequences $\mathcal{C} = (8, 16, 32, 64)$, $\mathcal{M} = (3, 4, 5, 6)$ and




Figure 6.7 – Normalized fairness of the configurations that Algorithm 1 returns when run for 100 throughput requirements S chosen randomly, for each N value.

$\mathcal{D} = (0, 1, 2, 3)$ ⁸. To evaluate fairness, we employ Jain’s fairness index $J(w)$ introduced in Section 5.2—where w is the time-horizon over which we evaluate J and it is expressed in number of PLC frames—and from now on, we evaluate this index at windows of N transmissions⁹, i.e., we use $J(N)$. First, we run simulations of all the possible configurations in $\{\mathcal{C} \times \mathcal{M} \times \mathcal{D}\}$, and we compute short-term fairness $J(N)$ and normalized throughput S . Let S_{min} be the minimum S achieved by all configurations in $\{\mathcal{C} \times \mathcal{M} \times \mathcal{D}\}$, and, similarly S_{max} be the maximum S . To test our algorithm, we draw 100 throughput requirements uniformly at random in $[S_{min}, S_{max}]$. Then, for each sample i with throughput S_i , we run Algorithm 1 that returns the configuration $config_i$.

Now, let J_i be the short-term fairness of the configuration $config_i$ at sample i , and let J_i^{max} denote the maximum short-term fairness of all configurations that satisfy the throughput constraint S_i . To evaluate the algorithm efficiency, we employ the *normalized fairness* which we define as J_i/J_i^{max} . The normalized fairness is a metric that evaluates the distance between the fairness of the configuration $config_i$ and the maximum achievable fairness, given the the throughput constraint S_i and the configuration search-space.

The results of the algorithm evaluation are presented in Figure 6.7. We present the normalized fairness of the configurations returned from 100 runs of Algorithm 1. We repeat the procedure described above for $2 \leq N \leq 10$. We observe that Algorithm 1 always returns a configuration with good fairness, given the throughput constraint¹⁰. Thus, it can be employed to optimize the performance for delay sensitive traffic operating with 1901. Because it uses the Drift model and binary search, the algorithm reduces the required search-space and it enables an efficient and fast method for determining good

⁸We use the factor $f = 2$, because the contention windows are also doubled between successive backoff stages. Algorithm 1 can be modified to include different values of f given the performance tradeoff of f in Table 5.1.

⁹We focus on the shortest possible horizon, see another example in Section 5.6.

¹⁰The minimum normalized fairness returned is about 0.8, but it was returned for very few throughput constraints S_i .

configurations in terms of jitter, without the need of running multiple time-consuming simulations.

6.5 Related Work and Discussion

The model introduced in this chapter is the first 1901 multi-user performance analysis that does not rely on the decoupling assumption. We adopted our analytical framework from similar 802.11 models. In this section, we review the 802.11 frameworks that either do not assume decoupling or validate the decoupling hypothesis¹¹.

Sharma et al. [57] study 802.11 without resorting to the decoupling assumption. They analyze an m -dimensional chain (m being the number of backoff stages) that describes the number of stations at each backoff stage, by assuming that the backoff counters follow a geometric distribution with the same mean as the real uniform one; this yields a constant transmission probability at any random slot for a station at a certain backoff stage. *Drift equations* capture the expected change of the number of stations at each backoff stage between two consecutive time slots, and their equilibrium point yields the average number of stations at each backoff stage in steady state. Similarly to Sharma et al. [57], we also use drift equations to obtain an accurate model for 1901. However, as the 1901 protocol is more complex than 802.11, so is our analysis; it differs substantially from the one of Sharma et al. [57]. The authors further show that the stochastic system converges to a deterministic one as $N \rightarrow \infty$. This convergence is proven under some necessary scaling assumptions that prevent the transition probabilities from converging to trivial values as $N \rightarrow \infty$. The resulting deterministic system lies in continuous time and it is a system of non-linear differential equations. Using the continuous-time approximation, the authors prove the global asymptotic stability of the equilibrium point for $m = 2$.

After, or in parallel with, the work of Sharma et al., there were many efforts aimed towards proving global asymptotic stability of 802.11 equilibrium points for general values of m . In particular, Cho et al. [54] prove the global asymptotic stability of an 802.11 system with K retries (where K here is equal to the number of backoff stages m), by employing asymptotic differential equations. The authors further explore the uniqueness and stability of heterogeneous networks, such as 802.11e, where contending stations might belong to classes with different configurations. Although they manage to prove 802.11e uniqueness of the equilibrium under some conditions, they state only the global-stability problem precisely without proving it. The complexity and the coupling introduced in the 802.11e challenge the study of the respective asymptotic differential equations. This situation is similar to the 1901 model, where the transmission probability of each station depends on the state of the whole system. In contrast, in simple 802.11, the transmission probability of any station is independent of the network state.

Bordenave et al. [56] prove the global stability of an 802.11 system under the conditions of infinite states ($m \rightarrow \infty$), binary exponential backoff, and a lower-bound for

¹¹Models of 802.11 and 1901 relying on the decoupling assumption are reviewed in Section 4.5.

the minimum contention window. Duffy [59] presents a thorough review of 802.11 works on the appropriateness of the decoupling approximation. He discusses related modeling problems under non-saturated queues per station and also under mesh topologies.

The veracity of the decoupling hypotheses is typically rigorously proven by *mean-field independence* with all works described above. In fact, the global asymptotic stability of the equilibrium point in the continuous-time approximation validates the decoupling assumption, as the stationary distribution of the number of stations in the stochastic system is concentrated at the equilibrium. The mathematical steps for proving the decoupling validity are typically the following: (i) Markov chain modeling of the system by relying on constant transition-probabilities assumptions for any station at a specific state; (ii) Scaling the system such that the transition probabilities do not converge to trivial values, such as 0 and 1, as $N \rightarrow \infty$ and proving convergence to a deterministic system in continuous time domain; (iii) Proving global asymptotic stability for the continuous-time dynamical system. In summary, we believe that the combination of the works by Benaïm and Le Boudec [58] and Cho et al. [54] can shed light on the mean-field independence of 1901. Steps (i) and (ii) above are already achieved by our work. Step (ii) is straightforward, given previous work [58] that established simple, mild conditions for the convergence of the stochastic system to the asymptotic differential equations. The most challenging step is clearly (iii), due to the high complexity of 1901. We leave this step for future work and we point to the work by Cho et al. [54] for the detailed methodology of proving global asymptotic stability for 802.11.

6.6 Summary

The decoupling assumption is commonly adopted for the analysis of MAC protocols, such as IEEE 802.11 and IEEE 1901. Building on our finding that there is strong coupling between 1901 stations, we propose the Drift model that does not rely on the decoupling assumption and, as a result, substantially improves the accuracy of previous 1901 analyses. The accuracy is particularly improved for networks with a small number of stations, which is the most frequent scenario in practice. Our analysis comprises performance in steady state, as well as in transient regime, and involves both the long-term and the short-term dynamics of 1901. We show that the Drift model admits a unique solution for a wide range of configurations.

We extensively compare the Drift and decoupling-assumption models. The differences in assumptions and complexity between the models enable different levels of accuracy and applications for improving performance. Given the different precision levels of the models, they can be combined to yield accurate throughput predictions for any number of stations N and any configuration. Both models have the same accuracy for large N , yet, the simplicity of decoupling-assumption enables throughput enhancements. The higher accuracy of the Drift model in the small- N regime can be exploited for applications that enhance other quality-of-service metrics, such as fairness and jitter. Notably, the Drift

model also analyzes the transient regime of the system; this enables us to evaluate the convergence time of various configurations and to study the global asymptotic stability of the system and the convergence from any initial state.

Finally, we explore the tradeoff between throughput and short-term fairness that exists in 1901 and that can be adjusted to accommodate different quality of service requirements in power-line networks. By exploiting the high accuracy of our Drift model and our previous study on this tradeoff, we present a heuristic algorithm for devising highly fair configurations that are based on the crucial features that a MAC protocol should satisfy, such as the reduction of aggressiveness when high traffic is present in the network. The algorithm efficiently regulates the tradeoff by returning the best configurations in terms of fairness, while satisfying a minimal throughput constraint.

7 Conclusion

In this dissertation, we have thoroughly investigated power-line communications and built the foundations for measuring, modeling and enhancing the performance of popular state-of-the-art solutions. The increasing connectivity demands and the interference or coverage problems in today's networks call for new communication technologies, such as PLC, that boost reliability, augment bandwidth and extend coverage. PLC is becoming increasingly popular, due to easy and high-throughput connectivity; indeed, we observe that it complements and often surpasses WiFi performance. Nevertheless, the research community has largely overlooked this technology, possibly due to its proprietary nature and the complexity of the protocols.

We have explored PLC from the PHY to MAC layers and enhanced both performance and delay. In Chapter 2, we have delved into the fundamental components of PLC and explained channel characteristics, implementations and the intricate protocols. We have also presented a thorough comparison between PLC and WiFi. In Chapter 3, we have measured the spatial and temporal variation of single PLC links. We find that PLC yields very good network-connectivity, but often with severe asymmetry of performance in one link. Furthermore, PLC capacity varies on three time-scales, and bad links exhibit the highest variation. By building on our spatio-temporal variation study, we have devised systematic link-metric guidelines for hybrid networks that include PLC. In Chapter 4, we have studied multi-user performance by proposing a simple model that relies on the so-called decoupling assumption. This assumption asserts that the backoff processes of the stations are independent. We have proven that the model yields a unique solution for a wide range of configurations and validated its accuracy by testbed experiments. By using this model, we have enhanced the multi-user throughput; our proposal consists in only modifying existing protocol-parameters, and we have corroborated the significant gains on a real testbed. In Chapter 5, we have investigated the short-term dynamics of multi-user scenarios and proven that the default PLC configuration is short-term unfair. We have shown that this unfairness introduces high jitter, crucial for delay-sensitive applications. Notably, we find that the high efficiency of the PLC MAC layer always comes at a cost of fairness; we have extensively investigated this compromise. PLC unfairness also weakens the accuracy of the decoupling assumption, under certain

configurations. To address this, in Chapter 6, we have introduced a new model that does not rely on the decoupling hypothesis: the dynamics of each station depends on the backoff processes of the other stations. By avoiding the decoupling assumption, our second model reaches high levels of accuracy, although it is more complex than the first one. We have extensively compared our two models and find that they offer different precision levels with different configurations and number of users. Therefore, we have employed our models for two different applications, that is, enhancing performance and tuning the tradeoff between throughput and fairness. With these applications, this thesis treats both best-effort and delay-sensitive applications.

We have enhanced the understanding of PLC performance and dynamics from the various aforementioned aspects. In addition, we have validated our findings with testbed experiments and have built an experimental framework for PLC hence, facilitating the future development and incorporation into hybrid networks. Yet, there are open questions about both PLC and hybrid networks.

Future networks need to simultaneously satisfy multiple stringent requirements from users, such as low power-consumption, full coverage, high reliability, rich bandwidth. Power consumption has not been explored in this thesis; it is interesting to compare WiFi and PLC with respect to this metric. Full coverage can be achieved by additional measurement studies of candidate technologies under various environments, such as residential, enterprise and public buildings. By exploiting models built on such measurements, hybrid-network topologies, that is, the type and location of stations, can be standardized for different environments. We have shown that PLC and WiFi temporal variations are diverse. Hence, by using both technologies, a temporal scheduler that has up-to-date capacity measurements and accurate capacity-predictions can augment network reliability. Achieving reliability is one of the most challenging problems in future hybrid networks. Hybrid algorithms need to take into account contention and efficiently share the available bandwidth among flows. A crucial challenge there is addressing delay requirements, especially in multi-hop topologies.

A very important research direction for future hybrid networks is the careful and reliable incorporation of low-rate technologies that aim at home automation and need to inter-operate with high-rate communications. Home-automation applications can be extremely delay-sensitive and, unfortunately, starved from bandwidth-hungry applications that utilize the same medium and spectrum. Future networks with both low- and high-rate technologies have to be scaled accordingly and to provide systematic topologies for seamless connectivity of all stations, potentially via multiple heterogeneous hops. It is clear that future networks need to handle diverse interference graphs, communications costs, and connectivity demands.

A Proofs

In this appendix, we give the proofs of all the mathematical formulations of Chapters 4–6. Table A.1 can be used as a quick reference for notation and equations. The definitions for x_k^i , bc_i , t_i , s_i are the same for our decoupling-assumption and Drift models, with the only difference that in the decoupling assumption one, we have $\gamma_i = \gamma$, for all $i \in \{0, m-1\}$. For both models, we introduce a new variable called B_i to facilitate the computations. B_i is defined in Lemma 4.

Notation	Value	Eq.
p_e	$\prod_{k=0}^{m-1} (1 - \tau_k)^{n_k}$	-
γ_i	$1 - \frac{p_e}{1 - \tau_i}$	(6.1)
x_k^i	$\sum_{j=d_i+1}^k \binom{k}{j} \gamma_i^j (1 - \gamma_i)^{k-j}$	(6.2)
t_i	$\sum_{k=d_i+1}^{CW_i-1} \frac{1}{CW_i} (1 - x_k^i) + \frac{d_i+1}{CW_i}$	(4.3)
s_i	$(1 - \gamma_i)t_i$	-
bc_i	$\frac{\sum_{k=d_i+1}^{CW_i-1} \left[(k+1)(1-x_k) + \sum_{j=d_i+1}^k j(x_j - x_{j-1}) \right]}{CW_i} + \frac{(d_i+1)(d_i+2)}{2CW_i}$	(6.3)
τ_i	$\frac{\sum_{k=d_i+1}^{CW_i-1} \frac{1}{CW_i} (1-x_k) + \frac{d_i+1}{CW_i}}{bc_i}$	(6.4)
B_i	$1/\tau_i - 1$	-
β_i	$\frac{\sum_{k=d_i+1}^{CW_i-1} \frac{1}{CW_i} \sum_{j=d_i+1}^k (x_j - x_{j-1})}{bc_i}$	(6.5)
K_i	$K_0 = 1, K_i = \frac{\tau_{i-1}\gamma_{i-1} + \beta_{i-1}}{\tau_i + \beta_i}, 1 \leq i \leq m-2, K_{m-1} = \frac{\tau_{m-2}\gamma_{m-2} + \beta_{m-2}}{\tau_{m-1}(1-\gamma_{m-1})}$	(6.6)
\hat{n}_i	$\hat{n}_0 = \frac{N}{\sum_{k=0}^{m-1} \prod_{j=0}^k K_j}, \hat{n}_i = \frac{N \prod_{j=0}^i K_j}{\sum_{k=0}^{m-1} \prod_{j=0}^k K_j} 1 \leq i \leq m-1$	(EQ)

Table A.1 – Summary of variables.

A.1 Decoupling-Assumption Model

The proof of Theorem 1 is built on a few lemmas and corollaries, which we state first.

Lemma 4. *Let γ be the collision probability. Let $B_i(\gamma)$ be the expected number of backoff slots between two transmissions attempts of a station that always remains at backoff stage i . Then, we have $B_i = bc_i/t_i - 1$, and B_i is an increasing function of γ , for any $0 \leq i \leq m - 1$.*

Proof. By its definition and by using the same rationale as the one employed to compute bc_i in (4.2), B_i is given recursively by

$$B_i = \frac{d_i(d_i + 1)}{2CW_i} + \sum_{j=d_i+1}^{CW_i-1} \frac{j(1 - x_j^i) + \sum_{k=d_i+1}^j (k + B_i)(x_k^i - x_{k-1}^i)}{CW_i}. \quad (\text{A.1})$$

Now, solving (A.1) over B_i , gives $B_i = bc_i/t_i - 1$, with bc_i and t_i given by (4.2) and (4.3).

To prove the second part of the lemma, we proceed as follows. (i) First, we compute $dB_i/d\gamma$. (ii) Second, we show that this derivative is positive at $\gamma = 1$. (iii) Third, we show that if the derivative is negative at some $0 < \gamma^* < 1$, it will also be negative at any value $\gamma > \gamma^*$. The proof then follows by contradiction: if the derivative was negative at some γ^* , it would also be negative at $\gamma = 1$, which would contradict (ii).

(i) After rearranging terms, (A.1) can be rewritten as

$$B_i = \frac{CW_i - 1}{2} + \frac{1}{CW_i} \sum_{j=d_i+1}^{CW_i-1} \left(B_i x_j^i - \sum_{k=d_i+1}^{j-1} x_k^i \right). \quad (\text{A.2})$$

The derivative of B_i can be computed as

$$\frac{dB_i}{d\gamma} = \sum_{k=d_i+1}^{CW_i-1} \frac{\partial B_i}{\partial x_k^i} \frac{dx_k^i}{d\gamma}.$$

The partial derivative $\partial B_i/\partial x_k^i$ can be computed from (A.2) as

$$\frac{\partial B_i}{\partial x_k^i} = \frac{B_i - (CW_i - 1 - k)}{CW_i} + \frac{\partial B_i}{\partial x_k^i} \sum_{j=d_i+1}^{CW_i-1} \frac{x_j^i}{CW_i},$$

which yields

$$\frac{dB_i}{d\gamma} = \frac{\sum_{k=d_i+1}^{CW_i-1} (B_i - (CW_i - 1 - k)) \frac{dx_k^i}{d\gamma}}{CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i}. \quad (\text{A.3})$$

To compute $dx_k^i/d\gamma$, we observe that x_k^i is the complementary cumulative function of a binomial distribution. By taking its derivative, we obtain

$$\frac{dx_k^i}{d\gamma} = \frac{k!}{(k-d_i-1)!d_i!} \gamma^{d_i} (1-\gamma)^{k-d_i-1}. \quad (\text{A.4})$$

(ii) Next, we show that $dB_i/d\gamma > 0$ at $\gamma = 1$. We have $x_k^i = 1$ at $\gamma = 1$ for all $d_i + 1 \leq k \leq CW_i - 1$. Given this and (A.1), we have

$$B_i = \frac{d_i(d_i+1)}{2CW_i} + \frac{CW_i - d_i - 1}{CW_i} (d_i + 1 + B_i). \quad (\text{A.5})$$

Solving (A.5) over B_i yields $B_i = CW_i - d_i/2 - 1$. Now, notice that $dx_k^i/d\gamma = 0$ at $\gamma = 1$ for all $d_i + 1 < k \leq CW_i - 1$, and $dx_{d_i+1}^i/d\gamma = d_i + 1$ from (A.4). Substituting in (A.3) yields $dB_i/d\gamma = d_i/2 + 1$, i.e., $dB_i/d\gamma > 0$.

(iii) Let us now assume that $dB_i/d\gamma < 0$ for some $\gamma^* < 1$. Let $l = \lceil CW_i - 1 - B_i(\gamma^*) \rceil$. Given (A.3), we can express $dB_i/d\gamma$ as the product of two terms, $dB_i/d\gamma = f_1(\gamma)f_2(\gamma)$, where

$$f_1(\gamma) \doteq \frac{dx_1^i/d\gamma}{CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i}, \quad f_2(\gamma) \doteq \sum_{k=d_i+1}^{CW_i-1} (B_i - (CW_i - 1 - k)) \frac{dx_k^i/d\gamma}{dx_1^i/d\gamma}.$$

We have $f_1(\gamma) > 0 \forall \gamma$, which implies $dB_i/d\gamma < 0$ if and only if $f_2(\gamma) < 0$. Also, we have

$$\begin{aligned} \frac{df_2(\gamma)}{d\gamma} &= \sum_{k=d_i+1}^{CW_i-1} \frac{dB_i}{d\gamma} \frac{dx_k^i/d\gamma}{dx_k^i/d\gamma} + \sum_{k=d_i+1}^{l-1} (B_i - (CW_i - 1 - k)) \frac{d}{d\gamma} \left(\frac{dx_k^i/d\gamma}{dx_1^i/d\gamma} \right) \\ &\quad + \sum_{k=l+1}^{CW_i-1} (B_i - (CW_i - 1 - k)) \frac{d}{d\gamma} \left(\frac{dx_k^i/d\gamma}{dx_1^i/d\gamma} \right), \end{aligned}$$

and

$$\frac{d}{d\gamma} \left(\frac{dx_k^i/d\gamma}{dx_1^i/d\gamma} \right) = -\frac{k!(l-d_i-1)!}{l!(k-d_i-1)!} (k-l)(1-\gamma)^{k-l-1},$$

which is positive for $k < l$ and negative for $k > l$. From the above equations, it follows that as long as $CW_i - 1 - (l-1) > B_i(\gamma) > CW_i - 1 - l$ and $dB_i/d\gamma < 0$, we have $df_2/d\gamma < 0$.

Building on the above, next we show that $B_i(\gamma)$ decreases for $\gamma \in [\gamma^*, \gamma^l]$, where γ^l is the γ value for which $B_i(\gamma^l) - (CW_i - 1 - l) = 0$. At $\gamma = \gamma^*$, we have $f_2(\gamma^*) < 0$, $dB_i/d\gamma < 0$ and $df_2/d\gamma < 0$. Let us assume that, before $B_i(\gamma)$ decreases down to $CW_i - 1 - l$, there is some $\hat{\gamma} > \gamma^*$ for which $dB_i/d\gamma \geq 0$. This implies that for some $\gamma' \in (\gamma^*, \hat{\gamma})$, $f_2(\gamma)$ has to stop decreasing, i.e., $df_2(\gamma')/d\gamma = 0$. Since $f_2(\gamma)$ decreases in $[\gamma^*, \gamma']$, we have $f_2(\gamma) < 0$ for $\gamma \in [\gamma^*, \gamma']$. Thus, $B_i(\gamma)$ decreases in $[\gamma^*, \gamma']$. As $CW_i - 1 - (l-1) > B_i(\gamma^*) > CW_i - 1 - l$ and (by assumption) $B_i(\gamma)$ does not reach $CW_i - 1 - l$, we also have $CW_i - 1 - (l-1) > B_i(\gamma') > CW_i - 1 - l$, which contradicts $df_2(\gamma')/d\gamma = 0$. Hence, our assumption does not hold, and $dB_i/d\gamma < 0$ until B_i reaches $CW_i - 1 - l$, i.e., $dB_i/d\gamma < 0$ for $\gamma \in [\gamma^*, \gamma^l]$.

Appendix A. Proofs

Following the same rationale for $\gamma \in [\gamma^l, \gamma^{l+1}]$, we can prove that $dB_i/d\gamma < 0$ for $\gamma \in [\gamma^l, \gamma^{l+1}]$. We can repeat this recursively to show that $dB_i/d\gamma < 0$ for $\gamma \in [\gamma^{l+1}, \gamma^{l+2}]$, $\gamma \in [\gamma^{l+2}, \gamma^{l+3}]$ until reaching $\gamma = 1$, which yields a contradiction because $dB_i/d\gamma > 0$ at $\gamma = 1$ from step (ii) above. \square

Corollary 1. *If $CW_{i+1} \geq 2CW_i - d_i - 1$, then we have $B_{i+1} > B_i$, for any $0 \leq i \leq m-2$.*

Proof. By Lemma 4, the minimum value of B_{i+1} is $B_{i+1}^{min} = (CW_{i+1} - 1)/2$ at $\gamma = 0$, and the maximum value of B_i is $B_i^{max} = CW_i - d_i/2 - 1$ at $\gamma = 1$. Setting $CW_{i+1} \geq 2CW_i - d_i - 1$, yields $B_{i+1}^{min} \geq B_i^{max}$, hence $B_{i+1} > B_i$ for all $\gamma \in [0, 1]$. \square

Corollary 2. *If $CW_{i+1} = CW_i$ and $d_{i+1} < d_i$, then we have $B_{i+1} \geq B_i$, for any $0 \leq i \leq m-2$.*

Proof. By the proof of Lemma 4, the equality holds for $\gamma = 0$, because $B_i(0) = (CW_i + 1)/2$. We now show that for $\gamma \in (0, 1]$, we have $B_{i+1} > B_i$.

If we prove $B_{i+1} > B_i$ when $CW_{i+1} = CW_i$, $d_{i+1} = d$ and $d_i = d + 1$, then the corollary follows by induction. Thus, we now show $B_{i+1} > B_i$ for this case, and we proceed as follows. Given (A.2), the difference $B_{i+1} - B_i$ can be computed as

$$B_{i+1} - B_i = \frac{\sum_{j=d+1}^{CW_i-1} (B_{i+1}x_j^{i+1} - \sum_{k=d+1}^{j-1} x_k^{i+1})}{CW_i} - \frac{\sum_{j=d+2}^{CW_i-1} (B_i x_j^i - \sum_{k=d+2}^{j-1} x_k^i)}{CW_i}. \quad (\text{A.6})$$

We have $x_j^{i+1} = x_j^i + \binom{j}{d+1} \gamma^{d+1} (1-\gamma)^{j-d-1}$. Let $\delta_j \doteq \binom{j}{d+1} \gamma^{d+1} (1-\gamma)^{j-d-1}$. Then, we have $\delta_j = x_j^{i+1} - x_j^i$. By rearranging the terms and solving over the difference $B_{i+1} - B_i$ in (A.6), and by using the definition of δ_j , we have

$$B_{i+1} - B_i = \frac{\sum_{k=d+1}^{CW_i-1} (B_{i+1} - (CW_i - 1 - k)) \delta_k}{CW_i - \sum_{j=d+2}^{CW_i-1} x_j^i}. \quad (\text{A.7})$$

By using (A.4), we now observe that

$$\delta_j = \frac{dx_j^{i+1}}{d\gamma} \frac{\gamma}{d+1}.$$

Substituting this in (A.7), yields

$$\begin{aligned} B_{i+1} - B_i &= \frac{\sum_{j=d+1}^{CW_i-1} (B_{i+1} - (CW_i - 1 - j)) \frac{\gamma}{d+1} \frac{dx_j^{i+1}}{d\gamma}}{CW_i - \sum_{j=d+2}^{CW_i-1} x_j^i} \\ &= \frac{CW_i - \sum_{j=d+1}^{CW_i-1} x_j^{i+1}}{CW_i - \sum_{j=d+2}^{CW_i-1} x_j^i} \frac{\gamma}{d+1} \left(\frac{dB_{i+1}}{d\gamma} \right) > 0, \end{aligned}$$

where the inequality holds by Lemma 4 and for all $\gamma \in (0, 1]$. \square

Corollary 3. *If $CW_{i+1} > CW_i$ and $d_{i+1} = d_i$, then we have $B_{i+1} > B_i$, for any $0 \leq i \leq m - 2$.*

Proof. If $B_{i+1} > B_i$ holds for $CW_{i+1} = CW_i + 1$, by using induction it is easy to see that it holds for any $CW_{i+1} > CW_i$. Thus, we prove the corollary for $CW_{i+1} = CW_i + 1$.

Because $d_i = d_{i+1}$ and by using (4.1), we have $x_k^i = x_k^{i+1}$ for all $d_i + 1 \leq k \leq CW_i - 1$. Given this, we have

$$\begin{aligned} B_{i+1} - B_i &= \frac{(CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i)^2}{(CW_i + 1 - \sum_{j=d_i+1}^{CW_i} x_j^i)(CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i)} \\ &\quad - \frac{(1 - x_{CW_i}^i) \frac{CW_i(CW_i-1)}{2}}{(CW_i + 1 - \sum_{j=d_i+1}^{CW_i} x_j^i)(CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i)} \\ &\quad + \frac{(1 - x_{CW_i}^i) \sum_{j=d_i+1}^{CW_i-1} (CW_i - 1 - j)x_j^i}{(CW_i + 1 - \sum_{j=d_i+1}^{CW_i} x_j^i)(CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i)}. \end{aligned}$$

By the definition of x_k^i in (4.1), we have $x_{CW_i}^i \geq x_k^i$, $d_i + 1 \leq k \leq CW_i - 1$, with equality at $\gamma = 0, 1$. This yields $\sum_{j=d_i+1}^{CW_i-1} x_j^i \leq (CW_i - d_i - 1)x_{CW_i}^i < CW_i x_{CW_i}^i$. We thus have

$$\begin{aligned} B_{i+1} - B_i &> \frac{CW_i(1 - x_{CW_i}^i)(CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i)}{(CW_i + 1 - \sum_{j=d_i+1}^{CW_i} x_j^i)(CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i)} \\ &\quad - \frac{(1 - x_{CW_i}^i) \frac{CW_i(CW_i-1)}{2}}{(CW_i + 1 - \sum_{j=d_i+1}^{CW_i} x_j^i)(CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i)} \\ &\quad + \frac{(1 - x_{CW_i}^i) \sum_{j=d_i+1}^{CW_i-1} (CW_i - 1 - j)x_j^i}{(CW_i + 1 - \sum_{j=d_i+1}^{CW_i} x_j^i)(CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i)} \\ &= \frac{(1 - x_{CW_i}^i) \left(\frac{CW_i(CW_i+1)}{2} - \sum_{j=d_i+1}^{CW_i-1} (j+1)x_j^i \right)}{(CW_i + 1 - \sum_{j=d_i+1}^{CW_i} x_j^i)(CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i)} \geq 0, \end{aligned}$$

where the last inequality holds because $x_j^i \leq 1$, for all i, j . □

We now present the proof of Lemmas 1 and 2, and Theorem 1, introduced in Section 4.2.

Proof of Lemma 1. Let X_i be the random variable denoting the number of slots that a station that starts in stage i spends in backoff before transmitting its current packet successfully. With this notation, it holds $\mathbb{E}[X] = \mathbb{E}[X_0]$. Let c_i and j_i denote the probabilities that a station at stage i ends this stage due to a collision, or due to sensing the medium busy $d_i + 1$ times, respectively. Note that $s_i + c_i + j_i = 1$ for all i . Additionally, let bc_{s_i} , bc_{c_i} and bc_{j_i} be the expected number of backoff slots that a station spends in backoff stage i , given that the station ends up redrawing its backoff counter due a

Appendix A. Proofs

packet successfully transmitted, due to a collision, or due to sensing the medium busy, respectively. From the law of total probability, we have

$$\mathbb{E}[X] = s_0bc_{s_0} + c_0bc_{c_0} + j_0bc_{j_0} + (c_0 + j_0)\mathbb{E}[X_1] = bc_0 + (1 - s_0)\mathbb{E}[X_1]. \quad (\text{A.8})$$

Repeating the above reasoning recursively for $\mathbb{E}[X_1], \mathbb{E}[X_2], \dots, \mathbb{E}[X_{m-1}]$, we have

$$\mathbb{E}[X_0] = bc_0 + (1 - s_0)\left(bc_1 + (1 - s_1)(bc_2 + \dots(1 - s_{m-2})\mathbb{E}[X_{m-1}])\right). \quad (\text{A.9})$$

Now, by applying the same reasoning as in (A.8), we have

$$\mathbb{E}[X_{m-1}] = bc_{m-1} + (1 - s_{m-1})\mathbb{E}[X_{m-1}].$$

Solving for $\mathbb{E}[X_{m-1}]$, we obtain $\mathbb{E}[X_{m-1}] = bc_{m-1}/s_{m-1}$. Plugging this expression into (A.9) concludes the proof. \square

Proof of Lemma 2. Similar to the proof of Lemma 1. \square

Proof of Theorem 1. By Brouwer's fixed-point theorem, since $\Gamma(G(\gamma))$ is a continuous function, there exists a fixed-point in $[0, 1]$. Furthermore, if $\Gamma(G(\gamma))$ is monotone, this fixed-point is unique. As $\Gamma(\tau) = 1 - (1 - \tau)^{N-1}$ is non-decreasing in γ , it is thus sufficient to show that $G(\gamma)$ is monotone in γ .

Let $Q(\gamma) = (1 - \gamma)\mathbb{E}[X]$. Then, we have $G(\gamma) = 1/Q(\gamma)$. Now, $G(\gamma)$ is non-increasing in γ if and only if $Q(\gamma)$ is non-decreasing in γ , which we show in the following.

We now use Lemma 4 to express $Q(\gamma)$ as a function of B_i . Replacing bc_i with $t_i(B_i + 1)$ in the expression for $\mathbb{E}[X]$ and using $s_i = (1 - \gamma)t_i$, $Q(\gamma)$ can be rewritten as

$$Q(\gamma) = \sum_{i=0}^{m-2} s_i(B_i + 1) \prod_{j=0}^{i-1} (1 - s_j) + \prod_{i=0}^{m-2} (1 - s_i)(B_{m-1} + 1).$$

The derivative of $Q(\gamma)$ with respect to γ is given by

$$\begin{aligned} \frac{dQ}{d\gamma} &= \sum_{i=0}^{m-2} s_i \frac{dB_i}{d\gamma} \prod_{j=0}^{i-1} (1 - s_j) + \prod_{i=0}^{m-2} (1 - s_i) \frac{dB_{m-1}}{d\gamma} \\ &\quad - \sum_{i=0}^{m-2} \frac{ds_i}{d\gamma} \left[\sum_{j=i+1}^{m-2} \left((B_j + 1) s_j \frac{\prod_{k=0}^{j-1} (1 - s_k)}{1 - s_i} \right) \right. \\ &\quad \left. - (B_i + 1) \prod_{j=0}^{i-1} (1 - s_j) + \frac{\prod_{j=0}^{m-2} (1 - s_j)}{1 - s_i} (B_{m-1} + 1) \right]. \end{aligned} \quad (\text{A.10})$$

From Lemma 4, we have that $dB_i/d\gamma > 0$. Thus, the first two terms in (A.10) are

positive and it follows that

$$\begin{aligned} \frac{dQ}{d\gamma} &> - \sum_{i=0}^{m-2} \frac{ds_i}{d\gamma} \left[\sum_{j=i+1}^{m-2} \left((B_j + 1) s_j \frac{\prod_{k=0}^{j-1} (1 - s_k)}{1 - s_i} \right) \right. \\ &\quad \left. - (B_i + 1) \prod_{j=0}^{i-1} (1 - s_j) + \frac{\prod_{j=0}^{m-2} (1 - s_j)}{1 - s_i} (B_{m-1} + 1) \right]. \end{aligned} \quad (\text{A.11})$$

In the proof of Lemma 4, we show that x_k^i is increasing with γ (see (A.4)). Thus, s_i is decreasing with γ , since, from (4.3), it holds

$$\frac{ds_i}{d\gamma} = -t_i - (1 - \gamma) \frac{\sum_{k=d_i+1}^{CW_i-1} dx_k^i/d\gamma}{CW_i} < 0.$$

From Corollaries 1–3 and Condition (4.5) of Theorem 1, $B_i(\gamma)$ is non-decreasing with i . Combining these two properties (i.e., $ds_i/d\gamma < 0$, $B_{i+1} \geq B_i$) for $i = m - 2$ with (A.11), we have

$$\begin{aligned} \frac{dQ}{d\gamma} &\geq - \sum_{i=0}^{m-3} \frac{ds_i}{d\gamma} \frac{1}{1 - s_i} \left[\sum_{j=i+1}^{m-2} \left((B_j + 1) s_j \prod_{k=0}^{j-1} (1 - s_k) \right) \right. \\ &\quad \left. - (B_i + 1) \prod_{j=0}^i (1 - s_j) + \prod_{j=0}^{m-2} (1 - s_j) (B_{m-1} + 1) \right]. \end{aligned}$$

By using $ds_i/d\gamma < 0$, $B_{i+1} \geq B_i$ in the above inequality and rearranging the factors in products involving the s_i 's, we have

$$\begin{aligned} \frac{dQ}{d\gamma} &\geq - \sum_{i=0}^{m-3} \frac{ds_i}{d\gamma} \frac{1}{1 - s_i} \left[(B_i + 1) \sum_{j=i+1}^{m-2} \left(s_j \prod_{k=0}^{j-1} (1 - s_k) \right) \right. \\ &\quad \left. - (B_i + 1) \prod_{j=0}^i (1 - s_j) + \prod_{j=0}^{m-2} (1 - s_j) (B_{m-1} + 1) \right] \\ &= - \sum_{i=0}^{m-3} \frac{ds_i}{d\gamma} \frac{\prod_{j=0}^{m-2} (1 - s_j)}{1 - s_i} (B_{m-1} - B_i) \geq 0, \end{aligned}$$

with equality at $\gamma = 0$. This completes the proof. □

A.2 Drift Model

The proof of uniqueness is based on a few lemmas and corollaries, which we state first. We employ Lemma 4 of the decoupling-assumption model, with the only difference that now B_i is a function of γ_i and not γ .

Corollary 4. τ_i is an decreasing function of γ_i , for any $0 \leq i \leq m - 1$.

Proof. From (6.4) and Lemma 4, it is easy to see that $\tau_i(\gamma_i) = 1/(B_i(\gamma_i) + 1)$. Thus, by Lemma 4, τ_i is an decreasing function of the collision probability γ_i . \square

Corollary 5. β_i is an increasing function of γ_i , for any $0 \leq i \leq m - 1$.

Proof. From (6.4) and (6.5) we have $\beta_i = 1/bc_i - \tau_i$. bc_i in (6.3) can be simplified as

$$bc_i = \frac{CW_i + 1}{2} - \frac{\sum_{k=d_i+1}^{CW_i-1} \sum_{j=d_i+1}^k x_j^i}{CW_i}. \quad (\text{A.12})$$

Since $dx_k^i/d\gamma_i > 0, k \geq d_i + 1$ (see (A.4)), bc_i is decreasing with γ_i . Hence, by using Corollary 4, β_i is increasing with γ_i . \square

Corollary 6. For any value of i , $dB_i/d\gamma_i < B_i^2/(1 - \gamma_i) \forall \gamma_i \in [0, 1]$, and $dB_i/d\gamma_i < B_i/\gamma_i$ for all $\gamma_i \in (0, 1]$.

Proof. We start with the first inequality. From (A.3), we have

$$\frac{dB_i}{d\gamma_i} < B_i \frac{1}{CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i} \sum_{k=d_i+1}^{CW_i-1} \frac{dx_k^i}{d\gamma_i}. \quad (\text{A.13})$$

To prove the lemma, we distinguish two cases: one for $d_i = 0$ and one for $d_i > 0$.

First, let us study (A.13) with $d_i = 0$. We have $dx_k^i/d\gamma_i = k(1 - \gamma_i)^{k-1}$ and $x_k^i = 1 - (1 - \gamma_i)^k$. Let $h(\gamma_i) = \sum_{k=0}^{CW_i-1} k(1 - \gamma_i)^k / \sum_{k=0}^{CW_i-1} (1 - \gamma_i)^k$. We now show that h decreases with γ_i . Let also $G(\gamma_i) = \sum_{k=0}^{CW_i-1} \gamma_i^k / \sum_{k=0}^{CW_i-1} k\gamma_i^k$. By Lemma 5.1 in [30], $G(\gamma_i)$ is strictly decreasing with γ_i in $[0, 1]$. Thus, $h(\gamma_i) = 1/G(1 - \gamma_i)$ is also strictly decreasing with γ_i in $[0, 1]$, and $h(\gamma_i) \leq h(0) = (CW_i - 1)/2$. Also, $B_i(\gamma_i) \geq (CW_i - 1)/2$ by Lemma 4. Given the above, we have $h(\gamma_i) \leq B_i(\gamma_i)$ and (A.13) yields

$$\frac{dB_i}{d\gamma_i} < \frac{B_i}{1 - \gamma_i} \frac{\sum_{k=0}^{CW_i-1} k(1 - \gamma_i)^k}{\sum_{k=0}^{CW_i-1} (1 - \gamma_i)^k} \leq \frac{B_i^2}{1 - \gamma_i}.$$

We now move to the case $d_i > 0$. From (A.4) and (4.1), we have $dx_k^i/d\gamma_i = k(x_k^i - x_{k-1}^i)/\gamma_i$. Thus, (A.13) yields

$$\frac{dB_i}{d\gamma_i} < \frac{B_i}{\gamma_i} \frac{CW_i x_{CW_i-1}^i - \sum_{k=d_i+1}^{CW_i-1} x_k^i}{CW_i - \sum_{k=d_i+1}^{CW_i-1} x_k^i} \leq \frac{B_i}{\gamma_i} x_{CW_i-1}^i. \quad (\text{A.14})$$

From the above, we have

$$\begin{aligned}
\frac{x_{CW_i-1}^i}{\gamma_i} &= \sum_{j=d_i+1}^{CW_i-1} \binom{CW_i-1}{j} \gamma_i^{j-1} (1-\gamma_i)^{CW_i-1-j} \\
&= \sum_{j=d_i+1}^{CW_i-1} \frac{CW_i-j}{j(1-\gamma_i)} \binom{CW_i-1}{j-1} \gamma_i^{j-1} (1-\gamma_i)^{CW_i-1-(j-1)} \\
&\leq \frac{CW_i-1}{2(1-\gamma_i)} \sum_{j=d_i+1}^{CW_i-1} \binom{CW_i-1}{j-1} \gamma_i^{j-1} (1-\gamma_i)^{CW_i-1-(j-1)} \\
&\leq \frac{CW_i-1}{2(1-\gamma_i)} \leq \frac{B_i}{1-\gamma_i}.
\end{aligned}$$

Thus, combining the above two equations we have $dB_i/d\gamma_i < B_i^2/(1-\gamma_i)$. Then, $dB_i/d\gamma_i < B_i/\gamma_i$ follows from (A.14), since $x_{CW_i-1}^i \leq 1$, which completes the proof. \square

Lemma 5. *Let us consider the expression of τ_i as a function of p_e resulting from combining (6.4) with (6.1). According to this expression, τ_i is an increasing function of p_e , for any $0 \leq i \leq m-1$.*

Proof. Since $\tau_i = 1/(B_i + 1)$, we need to show that $dB_i/dp_e < 0$. Note that

$$\frac{dB_i}{dp_e} = \frac{dB_i}{d\gamma_i} \frac{d\gamma_i}{dp_e}. \quad (\text{A.15})$$

From $\gamma_i = 1 - p_e/(1 - \tau_i) = 1 - p_e(B_i + 1)/B_i$, we have

$$\frac{d\gamma_i}{dp_e} = -\frac{B_i + 1}{B_i} + \frac{p_e}{B_i^2} \frac{dB_i}{dp_e}. \quad (\text{A.16})$$

Combining (A.15) and (A.16) yields

$$\frac{dB_i}{dp_e} = -\frac{dB_i}{d\gamma_i} \frac{B_i + 1}{B_i} \frac{1}{1 - \frac{p_e}{B_i^2} \frac{dB_i}{d\gamma_i}}. \quad (\text{A.17})$$

Let us distinguish two cases to prove this lemma, one for $\gamma_i = 1$ and the other for $0 \leq \gamma_i < 1$. First, for $\gamma_i = 1$, we have $p_e = 0$ from (6.1). From Lemma 4, we have $dB_i/d\gamma_i > 0$. Thus, (A.17) is smaller than 0.

We now look at the case $0 \leq \gamma_i < 1$. We have $dB_i/d\gamma_i > 0$ therefore, $dB_i/dp_e < 0$ as long as

$$\frac{dB_i}{d\gamma_i} < \frac{B_i^2}{p_e} = \frac{B_i(B_i + 1)}{1 - \gamma_i}. \quad (\text{A.18})$$

According to Corollary 6, $dB_i/d\gamma_i < B_i^2/(1 - \gamma_i)$, which is a sufficient condition for (A.18). This terminates the proof. \square

Appendix A. Proofs

The following lemmas relate to the stationary regime of (DYNSSYS). Let $\Phi(p_e) = \prod_{k=0}^{m-1} (1 - \tau_k(p_e))^{\hat{n}_k(p_e)}$. The following lemmas examine the function $\Phi(p_e)$, where each $\hat{n}_k(p_e)$ is a function of $\beta_i(p_e)$, $\gamma_i(p_e)$, $\tau_i(p_e)$, $0 \leq i \leq m-1$, as given by (EQ).

Lemma 6. *If (COND) is satisfied, then $\partial\Phi/\partial\beta_j > 0$, for any $0 \leq j < m-1$, and $\partial\Phi/\partial\beta_{m-1} = 0$.*

Proof. We consider the expression $\prod_{k=0}^{m-1} (1 - \tau_k)^{\hat{n}_k}$ as a function of τ_i , γ_i and β_i , where \hat{n}_i is computed as a function of τ_i , β_i and γ_i from (EQ). We show that if we increase β_j to β_j^* for a given j , and leave the remaining τ_i , γ_i and β_i values fixed, then $\prod_{k=0}^{m-1} (1 - \tau_k)^{\hat{n}_k}$ increases. First, we have $\partial\Phi/\partial\beta_{m-1} = 0$, because $\Phi(p_e)$ does not depend on β_{m-1} . We next study the cases with $0 \leq j < m-1$.

From (6.6), it can be seen that the new K_i values resulting from β_j^* , denoted by K_i^* , satisfy the following. If $j = 0$, then $K_1^* > K_1$ and $K_i^* = K_i$, $i > 1$ by (6.6). Thus, $\hat{n}_0^* < \hat{n}_0$ and $\hat{n}_i^* > \hat{n}_i$, $0 < i \leq m-1$. For $1 \leq j \leq m-2$, we have $\prod_{n=1}^i K_n^* = \prod_{n=1}^i K_n$, $i < j$ and $\prod_{n=1}^j K_n^* < \prod_{n=1}^j K_n$. We also have $\prod_{n=1}^i K_n^* > \prod_{n=1}^i K_n$, $i > j$, since

$$\frac{\prod_{n=1}^i K_n^*}{\prod_{n=1}^i K_n} = \frac{\tau_j \gamma_j + \beta_j^*}{\tau_j + \beta_j^*}, \quad \frac{\partial}{\partial \beta_j} \left(\frac{\tau_j \gamma_j + \beta_j}{\tau_j + \beta_j} \right) = \frac{\tau_j (1 - \gamma_j)}{(\tau_j + \beta_j)^2} > 0.$$

Let $\sigma = \sum_{i=1}^{m-1} \prod_{n=1}^i K_n$. We now show that $\sigma^* > \sigma$, i.e., $\partial\sigma/\partial\beta_j > 0$. For $j = m-2$, $\partial\sigma/\partial\beta_j > 0$ if and only if $\tau_{m-2}(1 - \gamma_{m-2}) - \tau_{m-1}(1 - \gamma_{m-1}) > 0$, which holds by (COND)¹.

For $j < m-2$, we have

$$\frac{\partial\sigma}{\partial\beta_j} = \prod_{l=1}^{j-1} \frac{K_l}{\tau_j + \beta_j} \left(-K_j + \frac{K_j \tau_j (1 - \gamma_j)}{\tau_{j+1} + \beta_{j+1}} \left(1 + \sum_{i=j+2}^{m-1} \prod_{n=j+2}^i K_n \right) \right). \quad (\text{A.19})$$

We prove $\partial\sigma/\partial\beta_j > 0$ by induction. We first show that $\partial\sigma/\partial\beta_j > 0$ for $j = m-3$, and then prove that if this holds for $j = k$, then it also holds for $j = k-1$. From (A.19), it can be seen that we need to show that:

$$\tau_j (1 - \gamma_j) \left(1 + \sum_{i=j+2}^{m-1} \prod_{n=j+2}^i K_n \right) - \tau_{i+1} - \beta_{i+1} > 0. \quad (\text{A.20})$$

For $j = m-3$, the above holds because of (COND). Now, assume that $\sigma^* > \sigma$ for $j = k$. Then, we show that $\sigma^* > \sigma$ holds also for $j = k-1$. Let us evaluate (A.20) at $j = k-1$. By using (A.20) as true for $j = k$, we have

$$\tau_{k-1} (1 - \gamma_{k-1}) \left(1 + \sum_{i=k+1}^{m-1} \prod_{n=k+1}^i K_n \right) > \tau_{k-1} (1 - \gamma_{k-1}) \left(1 + K_{k+1} \frac{\tau_{k+1} + \beta_{k+1}}{\tau_k (1 - \gamma_k)} \right).$$

¹According to (COND), τ_i decreases with i . From $\gamma_i = 1 - p_e/(1 - \tau_i)$, we have that γ_i increases with i .

Given (6.6) for K_{k+1} and (COND), the right-hand side of the above inequality yields

$$\tau_{k-1}(1 - \gamma_{k-1}) \left(1 + \frac{\tau_k \gamma_k + \beta_k}{\tau_k(1 - \gamma_k)} \right) > \tau_k(1 - \gamma_k) + \tau_k \gamma_k + \beta_k = \tau_k + \beta_k.$$

Thus, we have proven that if $\sigma^* > \sigma$ for $j = k$, then the same holds for $j = k - 1$. Given $\sigma^* > \sigma$ it follows that $\hat{n}_i^* < \hat{n}_i$ for $i \leq j$. Also, since $\hat{n}_i^* < \hat{n}_i$, $i \leq j$ and $\sum_k \hat{n}_k^* = \sum_k \hat{n}_k = N$, there must be some $l > j$ for which $\hat{n}_l^* > \hat{n}_l$. Since for $i \geq l + 1$, we have $\hat{n}_i = K_i \hat{n}_{i-1}$ with $K_i^* = K_i$, it holds that $\hat{n}_i^* > \hat{n}_i$, $i > l$. Thus,

$$\begin{aligned} \frac{\prod_{k=0}^{m-1} (1 - \tau_k)^{\hat{n}_k^*}}{\prod_{k=0}^{m-1} (1 - \tau_k)^{\hat{n}_k}} &= \prod_{k < l} (1 - \tau_k)^{\hat{n}_k^* - \hat{n}_k} \prod_{k \geq l} (1 - \tau_k)^{\hat{n}_k^* - \hat{n}_k} \\ &> (1 - \tau_l)^{\sum_{k < l} \hat{n}_k^* - \hat{n}_k} (1 - \tau_l)^{\sum_{k \geq l} \hat{n}_k^* - \hat{n}_k}, \end{aligned}$$

because of (COND). As $\sum_k \hat{n}_k^* = \sum_k \hat{n}_k = N$, the above is larger than 1, which proves the lemma. \square

Lemma 7. *If (COND) is satisfied, then $\partial\Phi/\partial\gamma_j > 0$, for any $0 \leq j \leq m - 1$.*

Proof. The proof is similar to the one of Lemma 6. It can be easily seen from (6.6) that if γ_j increases to γ_j^* , we have

$$\prod_{n=1}^i K_n^* = \prod_{n=1}^i K_n, \quad i \leq j \quad \text{and} \quad \prod_{n=1}^i K_n^* > \prod_{n=1}^i K_n, \quad i > j.$$

Note that the above holds for $0 \leq j \leq m - 2$. For $j = m - 1$ we have

$$\prod_{n=1}^i K_n^* = \prod_{n=1}^i K_n, \quad i < m - 1, \quad \prod_{n=1}^i K_n^* > \prod_{n=1}^i K_n, \quad i = m - 1.$$

Thus, as $\sigma^* > \sigma$ (with $\sigma = \sum_{i=1}^{m-1} \prod_{n=1}^i K_n$) also holds here, it is $\hat{n}_i^* < \hat{n}_i$ for $i \leq j$ and $\hat{n}_i^* > \hat{n}_i$ for $i > j$, with $0 \leq j \leq m - 2$. For $j = m - 1$ we have $\hat{n}_i^* < \hat{n}_i$ for $i < m - 1$ and $\hat{n}_i^* > \hat{n}_i$ for $i = m - 1$. Then, following the same reasoning as for the previous lemma, it can be seen that $\prod_{k=0}^{m-1} (1 - \tau_k)^{\hat{n}_k^*} > \prod_{k=0}^{m-1} (1 - \tau_k)^{\hat{n}_k}$, which proves the lemma. \square

Lemma 8. *If (COND) is satisfied, then $\partial\Phi/\partial\tau_j < 0$, for any $0 \leq j \leq m - 1$.*

Proof. When τ_j increases to τ_j^* , $\prod_{n=1}^i K_n^* = \prod_{n=1}^i K_n$ for $i < j$, and $\prod_{n=1}^i K_n^* < \prod_{n=1}^i K_n$ for $i = j$. For $i > j$ we have

$$\prod_{n=1}^i K_n^* = \prod_{n=1}^i K_n \frac{\tau_j^* \gamma_j + \beta_j}{\tau_j^* + \beta_j} \frac{\tau_j + \beta_j}{\tau_j \gamma_j + \beta_j}, \quad \text{and} \quad \frac{\partial}{\partial \tau_j} \left(\frac{\tau_j \gamma_j + \beta_j}{\tau_j + \beta_j} \right) < 0.$$

Appendix A. Proofs

Thus, $\prod_{n=1}^i K_n^* < \prod_{n=1}^i K_n$ for $i \geq j$. These yield $\sigma^* < \sigma$ (with $\sigma = \sum_{i=1}^{m-1} \prod_{n=1}^i K_n$). From the above, $\hat{n}_i^* > \hat{n}_i$ for $i < j$ and $\hat{n}_j^* < \hat{n}_j$, because $K_j^* < K_j$. Hence, by using similar arguments as in the two previous lemmas, it follows that $\prod_{k \neq j} (1 - \tau_k)^{\hat{n}_k}$ decreases. If we show that $(1 - \tau_j)^{\hat{n}_j}$ also decreases, the lemma will be proven. Note that

$$\frac{\partial(1 - \tau_j)^{\hat{n}_j}}{\partial \tau_j} = -\hat{n}_j(1 - \tau_j)^{\hat{n}_j-1} + \ln(1 - \tau_j) \frac{\partial \hat{n}_j}{\partial \tau_j} (1 - \tau_j)^{\hat{n}_j}. \quad (\text{A.21})$$

Computing the partial derivative of \hat{n}_j , we have

$$\frac{\partial \hat{n}_j}{\partial \tau_j} = -\frac{\hat{n}_j}{\tau_j + \beta_j} + \frac{\partial \hat{n}_{j-1}}{\partial \tau_j} K_j \geq -\frac{\hat{n}_j}{\tau_j + \beta_j} \geq -\frac{\hat{n}_j}{\tau_j}, \quad (\text{A.22})$$

because $\partial \hat{n}_{j-1} / \partial \tau_j > 0$ from above. This inequality holds for any $0 < j \leq m - 1$. By taking into account that $n_0 = N/\sigma$, it is easy to see that it also holds for $j = 0$. Note that, in this case, $\partial \sigma / \partial \tau_0 > 0$. For this case, we have

$$\frac{\partial \hat{n}_0}{\partial \tau_0} = -\frac{\hat{n}_0 (\sigma - 1) \gamma_0}{\sigma \tau_0 \gamma_0 + \beta_0} \geq -\frac{\hat{n}_0 \gamma_0}{\tau_0 \gamma_0 + \beta_0} \geq -\frac{\hat{n}_0}{\tau_0}.$$

Combining (A.21) and (A.22) yields

$$\frac{\partial(1 - \tau_j)^{\hat{n}_j}}{\partial \tau_j} \leq \frac{\hat{n}_j(1 - \tau_j)^{\hat{n}_j}}{\tau_j} \left(-\frac{\tau_j}{1 - \tau_j} - \ln(1 - \tau_j) \right).$$

As $-x/(1-x) < \ln(1-x)$, the above is smaller than 0, which proves the lemma. \square

Proof of Theorem 2. Recall that $p_e = \prod_{k=0}^{m-1} (1 - \tau_k)^{\hat{n}_k}$. For any value of p_e , τ_i can be computed from the fixed-point equation that results from combining (6.1) (i.e., $\gamma_i = 1 - p_e / (1 - \tau_i)$) with (6.4), where (6.4) is expressed as a function of γ_i through (4.1). Hence, τ_i can be computed as a function of p_e , and so can γ_i , and β_i . Now, \hat{n}_i can also be computed as a function of p_e using (EQ). Then, a solution of (EQ) has to satisfy the following equation:

$$p_e = \Phi(p_e). \quad (\text{A.23})$$

It can be seen that (A.23) has at least one fixed-point. $\Phi(p_e)$ is defined in $[0, 1 - \tau_{max}]$, where $\tau_{max} \doteq 2/(CW_0 + 1)$ is the maximum transmission probability at stage 0. Observe that $\Phi(0) > 0$ and $\Phi(1 - \tau_{max}) < 1 - \tau_{max}$ thus, by the intermediate value theorem, $\Phi(p_e)$ has at least one fixed-point in $[0, 1 - \tau_{max}]$. We now show that (A.23) has only one fixed-point. To this end, we show that $\Phi(p_e)$ is monotonically decreasing with p_e . The derivative of $\Phi(p_e)$ can be written as

$$\frac{d\Phi(p_e)}{dp_e} = \sum_{j=0}^{m-1} \left(\frac{\partial \Phi}{\partial \gamma_j} \frac{d\gamma_j}{dp_e} + \frac{\partial \Phi}{\partial \beta_j} \frac{d\beta_j}{dp_e} + \frac{\partial \Phi}{\partial \tau_j} \frac{d\tau_j}{dp_e} \right). \quad (\text{A.24})$$

We now examine separately each of the partial derivative products of (A.24) with respect to γ_j , β_j and τ_j . To prove the theorem, we rely on our analysis above. First, Lemmas 4 and 5 imply respectively that $d\gamma_j/d\tau_j < 0$ and $d\tau_j/dp_e < 0$. Because $d\gamma_j/dp_e = (d\gamma_j/d\tau_j) \cdot (d\tau_j/dp_e)$, we have $d\gamma_j/dp_e < 0$. Also, from Lemma 7, we have $\partial\Phi/\partial\gamma_j > 0$. Thus, the first product of partial derivatives in (A.24) is negative for all j . Second, from Lemma 6, we have $\partial\Phi/\partial\beta_j \geq 0$. Now, Corollary 5 states that $d\beta_j/d\gamma_j > 0$ and we have shown above that $d\gamma_j/dp_e < 0$. Hence, we have $d\beta_j/dp_e < 0$ thus, the second product of partial derivatives in (A.24) is also negative. Third, from Lemma 8 we have $\partial\Phi/\partial\tau_j < 0$, and from Lemma 5 we have $d\tau_j/dp_e > 0$. We have shown that all the partial derivative products of (A.24) are negative, so $\Phi(p_e)$ is monotonically decreasing with p_e .

Since (A.24) is strictly negative and (A.23) admits at least one fixed-point, there exists a unique value for p_e that solves (A.23). Computing the corresponding value for γ_i by (6.1), we have a solution to (EQ). The uniqueness of the solution then follows from the fact that all relationships between τ_i , β_i , γ_i and p_e are bijective, and any solution must satisfy (A.23), which (as we have shown) has only one solution. \square

Proof of Lemma 3. The proof goes by contradiction. Let us assume that there exists a solution \mathbf{n}^s , such that the corresponding values of τ_i^s , γ_i^s , τ_{i+1}^s and γ_{i+1}^s, p_e^s satisfy $\tau_i^s(\gamma_i^s) < \tau_{i+1}^s(\gamma_{i+1}^s)$ and (consequently) $\gamma_i^s > \gamma_{i+1}^s$. Note that, for any n_i distribution, γ_i and γ_{i+1} satisfy (6.1). Due to (6.1), we have

$$\frac{1 - \gamma_i^s}{1 - \gamma_{i+1}^s} = \frac{1 - \tau_{i+1}^s}{1 - \tau_i^s}. \quad (\text{A.25})$$

Let us fix τ_i and γ_i to the values given by the solution described above, and vary γ_{i+1} by choosing different values of ν , defined as $1 - \gamma_{i+1} \doteq \nu(1 - \gamma_i^s)$. For each ν , we first compute τ_{i+1} that corresponds to this γ_{i+1} , and then we compute the expression $(1 - \tau_i^s)/(1 - \tau_{i+1}^s)$ that results from this τ_{i+1} and the (fixed) τ_i^s value. Then, if such a solution s exists, there must be some value of $\nu \geq 1$ for which

$$r(\nu) \doteq \frac{1 - \tau_i^s}{1 - \tau_{i+1}^s(\nu)} = \nu,$$

because of (A.25) and the definition of ν . Next, we show that such a ν does not exist, which contradicts our initial assumption.

By hypothesis, for $\nu = 1$ we have $\gamma_{i+1} = \gamma_i^s$, so $\tau_i(\gamma_i^s) > \tau_{i+1}(\gamma_i^s)$, and $r(1) < 1$. A sufficient condition to ensure that there exists no $\nu > 1$ value for which $r(\nu) = \nu$ is that the derivative of $r(\nu)$, i.e. $dr(\nu)/d\nu$, does not exceed 1 in the region $\nu \geq 1$. To prove this,

Appendix A. Proofs

we proceed as follows.

$$\begin{aligned}
\frac{dr(\nu)}{d\nu} &= \frac{1 - \tau_i^s}{(1 - \tau_{i+1})^2} \frac{d\tau_{i+1}}{d\nu} = \frac{1 - \tau_i^s}{(1 - \tau_{i+1})^2} \frac{d\tau_{i+1}}{d\gamma_{i+1}} \frac{d\gamma_{i+1}}{d\nu} \\
&= -\frac{(1 - \tau_i^s)(1 - \gamma_i^s)}{(1 - \tau_{i+1})^2} \frac{d\tau_{i+1}}{dB_{i+1}} \frac{dB_{i+1}}{d\gamma_{i+1}} \\
&= \frac{(1 - \tau_i^s)(1 - \gamma_i^s)\tau_{i+1}^2}{(1 - \tau_{i+1})^2} \frac{dB_{i+1}}{d\gamma_{i+1}} \\
&\stackrel{\nu \geq 1}{\leq} \frac{(1 - \gamma_{i+1})\tau_{i+1}^2}{(1 - \tau_{i+1})^2} \frac{dB_{i+1}}{d\gamma_{i+1}} = \frac{(1 - \gamma_{i+1})}{B_{i+1}^2} \frac{dB_{i+1}}{d\gamma_{i+1}}.
\end{aligned}$$

From the above, it is sufficient to prove $dB_{i+1}/d\gamma_{i+1} < B_{i+1}^2/(1 - \gamma_{i+1})$. This is shown in Corollary 6 in Appendix for $\gamma_{i+1} \in [0, 1)$. For $\gamma_{i+1} = 1$, we also have $\gamma_i = 1$ by (6.1). Thus, a solution s cannot exist and $\tau_i(\gamma_i) > \tau_{i+1}(\gamma_{i+1})$. \square

Proof of Theorem 3. We analyze two cases: 1) $d_{i+1} = d_i$; 2) $d_{i+1} \neq d_i$.

1) We start for the case $d_{i+1} = d_i$. By using Lemma 3, we need only to prove that $\tau_{i+1}(\gamma_i) < \tau_i(\gamma_i)$. This is true by Corollary 3, given that $\tau_i(\gamma_i) = 1/(B_i(\gamma_i) + 1)$. 2) We now look at the case $d_{i+1} \neq d_i$. The result for this case follows from Corollary 1 and given that $\tau_i(\gamma_i) = 1/(B_i(\gamma_i) + 1)$. \square

Proof of Theorem 4. The proof follows from the work by Benaïm and Le Boudec[58]; according to this paper, it is sufficient to verify that the system analyzed satisfies the five assumptions given in [58], referred to as H1–H5. As our system does not have the so-called “common resource” (an additional entity with which the stations can interact), assumptions H1 and H4 are not applicable here. By taking the “vanishing intensity” $\epsilon(N) = 1/N$, it follows that $\lim_{N \rightarrow \infty} \mathbf{F}(\mathbf{y})/\epsilon(N)$ can be expressed as a function of $\tau_i(\rho)$ and $\beta_i(\rho)$ hence, assumption H2 (“*Intensity vanishes at a rate $\epsilon(N)$* ”) is satisfied.

To validate assumption H3 (“*Second moment of number of object transitions per time slot has an upper bound of order $N^2\epsilon^2(N)$* ”), we proceed as follows. Let $\alpha = \max_{i,p_i}(\tau_i(1 - p_i), \tau_i p_i + \beta_i)$. Note that $\alpha < 1^2$. Then, the probability that a station changes its state is upper bounded by α/N , and the number of stations that change their state is stochastically upper bounded by a random variable W_N that follows the binomial distribution $\text{Bin}(N, \alpha/N)^3$. Thus, we have $\mathbb{E}[W_N^2] = \mathbb{E}[W_N]^2 + \text{Var}[W_N] < \alpha^2 + \alpha$, which implies that assumption H3 is satisfied.

Finally, both τ_i and β_i are smooth functions of γ_i that in turn is a smooth function of the n_i 's, because all the aforementioned functions are continuous and continuously differentiable. Hence, the transition probabilities are smooth functions of the y_i 's. In addition, the transition probabilities chosen are also smooth functions of N . This

²We have $\alpha < 1$, given that $\tau_i(1 - p_i) < 1$ and $\tau_i p_i + \beta_i < 1$ (the latter follows from $\tau_i + \beta_i = 1/bc_i < 1$).

³Note that, under our network model, a station changes its state independently of the transitions of the other stations.

also holds for the boundaries of the transition probabilities, including the case when $N \rightarrow \infty$. Therefore, assumption H5 (“ $\mathbf{F}(\mathbf{n})$ is a smooth function of $1/N$ and \mathbf{n} ”) is also satisfied. \square

Proof of Theorem 5. Let $\kappa(t) = \tau_0 y_0(t) + \tau_1 y_1(t)$. Then, $\rho = 1 - e^{-\kappa}$ and (ODE) for $m = 2$ is given by

$$\begin{aligned}\frac{dy_0(t)}{dt} &= -y_0(\tau_0(1 - e^{-\kappa}) + \beta_0) + y_1\tau_1 e^{-\kappa} \\ \frac{dy_1(t)}{dt} &= y_0(\tau_0(1 - e^{-\kappa}) + \beta_0) - y_1\tau_1 e^{-\kappa},\end{aligned}$$

where the τ_i 's and β_i 's are functions of ρ , or equivalently of κ .

Let us consider the Lyapunov function $L(\mathbf{y}) = (y_0(t) - \hat{y}_0)^2 + (y_1(t) - \hat{y}_1)^2$, where \hat{y}_0 and \hat{y}_1 are the values of $y_0(t)$ and $y_1(t)$ at the equilibrium point. If at some time t we have $y_0(t) > \hat{y}_0$, this implies $y_1(t) < \hat{y}_1$ (since $y_0(t) + y_1(t) = 1$) and $\kappa(t) > \hat{\kappa}$. The latter can be seen by contradiction. Let us assume $\kappa(t) < \hat{\kappa}$. Then, from $\kappa = -\ln(1 - \rho)$ and Corollary 4 we have $\tau_0 > \hat{\tau}_0$ and $\tau_1 > \hat{\tau}_1$. Thus, $\kappa = \tau_0 y_0(t) + \tau_1 y_1(t) > \hat{\tau}_0 y_0(t) + \hat{\tau}_1 y_1(t) = y_0(\hat{\tau}_0 - \hat{\tau}_1) + \hat{\tau}_1 > \hat{y}_0(\hat{\tau}_0 - \hat{\tau}_1) + \hat{\tau}_1 = \hat{\tau}_0 \hat{y}_0(t) + \hat{\tau}_1 \hat{y}_1(t) = \hat{\kappa}$, since we have $\hat{\tau}_0 > \hat{\tau}_1$ from (COND). This contradicts the initial assumption.

Given $\kappa(t) > \hat{\kappa}$, we have $\tau_0(1 - e^{-\kappa}) + \beta_0 > \hat{\tau}_0(1 - e^{-\hat{\kappa}}) + \hat{\beta}_0$. This can be seen as follows. By employing a similar reasoning to Corollary 5, we have $\partial(\tau_0 + \beta_0)/\partial\kappa > 0$. We also have $-\partial(\tau_0 e^{-\kappa})/\partial\kappa > 0$ from Lemma 4. Then, adding both expressions we obtain $\partial(\tau_0(1 - e^{-\kappa}) + \beta_0)/\partial\kappa > 0$. Given $\kappa(t) > \hat{\kappa}$, we also have $\tau_1 e^{-\kappa} < \hat{\tau}_1 e^{-\hat{\kappa}}$. Thus,

$$\begin{aligned}\frac{dy_0(t)}{dt} &= -y_0(\tau_0(1 - e^{-\kappa}) + \beta_0) + y_1\tau_1 e^{-\kappa} \\ &< -y_0(\hat{\tau}_0(1 - e^{-\hat{\kappa}}) + \hat{\beta}_0) + y_1\hat{\tau}_1 e^{-\hat{\kappa}} \\ &< -\hat{y}_0(\hat{\tau}_0(1 - e^{-\hat{\kappa}}) + \hat{\beta}_0) + \hat{y}_1\hat{\tau}_1 e^{-\hat{\kappa}} = 0.\end{aligned}$$

Since $dy_0(t)/dt + dy_1(t)/dt = 0$, this in turn implies $dy_1(t)/dt > 0$. Putting all this together yields

$$\frac{dL(\mathbf{y})}{dt} = 2(y_0(t) - \hat{y}_0)\frac{dy_0(t)}{dt} + 2(y_1(t) - \hat{y}_1)\frac{dy_1(t)}{dt} < 0.$$

Following a similar reasoning, it can be seen that if $y_0(t) < \hat{y}_0$, then $dy_0(t)/dt > 0$ and $dy_1(t)/dt < 0$. As a consequence, we have $dL(\mathbf{y})/dt < 0$ also in this case. Therefore, the system is globally asymptotically stable. \square

A.3 Distribution of Inter-transmissions

We next present the computation of the distribution of inter-transmissions for IEEE 1901, given in Proposition 1 of Chapter 5. We assume that the backoff-counters are continuous random variables, thus there are no collisions between the stations. The station with the smallest backoff counter wins the channel whenever the stations resume or choose their backoff counters. Thus, the distribution of the number of inter-transmissions K , $\mathbb{P}(K = k)$, is computed by integrating the probability density functions of M backoff counters over the appropriate region of \mathbb{R}^M (M depends on k as shown in the next paragraphs). We next prove the results for CA0/CA1 priorities (see Table 2.2).

We assume that we have two stations A and B and that both stations start at the same state, i.e., backoff stage 0. We tag one of the stations in order to compute its distribution of K . Let B be the tagged station and A be the non-tagged station. We denote by $b^{(i)}$ the backoff counter that B draws when it enters backoff stage i . Thus, $b^{(0)}, b^{(1)}, b^{(2)}, b^{(3)}$ are uniformly distributed over the ranges $[0, 8], [0, 16], [0, 32], [0, 64]$, respectively. After reaching the maximum backoff stage which is 3, the station does not change the contention-window value. Thus, $b^{(i)}$ are uniformly distributed over the range $[0, 64]$ for $i \geq 3$, and B chooses a new backoff counter $b^{(i)}$ after every 16 consecutive transmissions of A , due to the deferral-counter value at this backoff stage (which is 15). We denote by a_k the backoff counter of A before the k th consecutive transmission, and a_k is uniformly distributed over the range $[0, 8]$, for all $k \geq 0$.

Due to the definition of 1901 protocol and to the continuous-time backoff counters assumption, stations change a backoff stage only when they sense the medium busy and their deferral counter is 0. As A is the one that transmits successfully before B wins the channel, A is always at backoff stage 0. B senses the medium busy and depending on the deferral-counter value after k consecutive transmissions of A , B chooses a new backoff counter. All a_k are mutually independent and independent of all $b^{(i)}$, because in the real protocol stations choose backoff counters independently of the other stations and of the backoff counters the stations selected in the past.

We normalize the distributions of the backoff counters, so that all a_k and $b^{(0)}$ are uniformly distributed over the range $[0, 1]$. Hence, $b^{(1)}, b^{(2)}, b^{(3)}$ are uniformly distributed over the ranges $[0, 2], [0, 4], [0, 8]$, respectively. This shows that the distribution of K does not depend on the minimum contention window value as long as we keep the binary exponential backoff of class CA1 of 1901. We denote by DC_B and BC_B the random variables that represent the backoff- and deferral-counter values at any state k , as shown in Table A.2. Given the above, $\mathbb{P}(K = k)$ is computed as follows:

- For $k = 0$, B transmits, thus $\mathbb{P}(K = 0) = P(a_1 \geq b^{(0)})$.
- For $k = 1$, A transmits once and then B transmits. Because DC_B is 0 when B enters backoff stage 0 (see Table 2.2), B enters backoff stage 1 and doubles CW

A.3. Distribution of Inter-transmissions

k	DC_B	BC_B
0	0	$b^{(0)}$
1	0	$b^{(0)}$
2	1	$b^{(1)}$
3	0	$b^{(1)} - a_2$
4	3	$b^{(2)}$
5	2	$b^{(2)} - a_4$
6	1	$b^{(2)} - a_4 - a_5$
7	0	$b^{(2)} - a_4 - a_5 - a_6$
8	15	$b^{(3)}$
9	14	$b^{(3)} - a_8$
\vdots	\vdots	\vdots
23	0	$b^{(3)} - a_8 - \sum_{i=8}^{22} a_i$
24	15	$b^{(4)}$

Table A.2 – The deferral counter of B , DC_B , and the backoff counter of B , BC_B , before the k th transmission of A .

after sensing the medium busy once. Hence, we have

$$\mathbb{P}(K = 1) = \mathbb{P}\left(a_1 < b^{(0)} \text{ and } a_2 \geq b^{(1)}\right).$$

- For $k = 2$, A transmits twice and then B transmits, thus

$$\mathbb{P}(K = 2) = \mathbb{P}\left(a_1 < b^{(0)} \text{ and } a_2 < b^{(1)} \text{ and } a_3 \geq b^{(1)} - a_2\right),$$

because DC_B is 1 when B enters backoff stage 1.

- For $k = 3$, A transmits three times and then B transmits:

$$\mathbb{P}(K = 3) = \mathbb{P}\left(a_1 < b^{(0)} \text{ and } a_3 < b^{(1)} - a_2 \text{ and } a_4 \geq b^{(2)}\right),$$

because DC_B becomes 0 after one transmission of A , at backoff stage 1.

- For $k \in [4, 6]$, A transmits k times and then B transmits:

$$\begin{aligned} \mathbb{P}(K = k) &= \mathbb{P}\left(a_1 < b^{(0)}\right) \mathbb{P}\left(a_3 + a_2 < b^{(1)}\right) \\ &\quad \cdot \mathbb{P}\left(a_k < b^{(2)} - \sum_{i=4}^{k-1} a_i \text{ and } a_{k+1} \geq b^{(2)} - \sum_{i=4}^k a_i\right), \end{aligned}$$

due to the independence of the backoff counters. After the k th transmission of A , B 's backoff counter equals $b^{(2)} - \sum_{i=4}^k a_i$.

- When $k = 7$, DC_B is 0. Thus, at the 8th attempt of A , B has a backoff counter

Appendix A. Proofs

$b^{(3)}$ and $DC_B = 15$. Thus,

$$\mathbb{P}(K = 7) = \mathbb{P}\left(a_1 < b^{(0)} \text{ and } a_3 + a_2 < b^{(1)}\right) \mathbb{P}\left(a_7 < b^{(2)} - \sum_{i=4}^6 a_i \text{ and } a_8 \geq b^{(3)}\right).$$

- For $8 \leq k \leq 22$, B decrements DC_B by 1 after each successful transmission of A , and the backoff counter is $b^{(3)} - \sum_{i=8}^k a_i$ when A has transmitted k times. Hence,

$$\begin{aligned} \mathbb{P}(K = k) &= \mathbb{P}\left(a_1 < b^{(0)}\right) P\left(a_3 + a_2 < b^{(1)}\right) \mathbb{P}\left(a_7 < b^{(2)} - \sum_{i=4}^6 a_i\right) \\ &\quad \cdot \mathbb{P}\left(a_k < b^{(3)} - \sum_{i=8}^{k-1} a_i \text{ and } a_{k+1} \geq b^{(3)} - \sum_{i=8}^k a_i\right). \end{aligned}$$

- When $k = 23$, $DC_B = 0$, thus at the next attempt B chooses a new backoff counter $b^{(4)}$ and the new DC_B is 15. Therefore, we have

$$\begin{aligned} \mathbb{P}(K = 23) &= \mathbb{P}\left(a_1 < b^{(0)}\right) \mathbb{P}\left(a_3 + a_2 < b^{(1)}\right) \\ &\quad \cdot \mathbb{P}\left(a_7 < b^{(2)} - \sum_{i=4}^6 a_i\right) \mathbb{P}\left(a_{23} < b^{(3)} - \sum_{i=8}^{22} a_i\right) \mathbb{P}\left(a_{24} \geq b^{(4)}\right). \end{aligned}$$

We now explain the computation details of the above probabilities. For $k < 15$ we integrate first with respect to the appropriate $b^{(i)}$, because the maximum value of the sum of the backoff counters of A is less or equal than the maximum value of $b^{(i)}$. For $15 \leq k \leq 22$ the maximum value of the sum of the backoff counters of A is larger than the maximum value of $b^{(i)}$. Thus, we define a new variable which is the sum of the backoff counters of A and we integrate first with respect to this random variable. Let U be the sum of n independent random variables uniformly distributed over the range $[0, 1]$. Then, U follows the Irwin-Hall distribution with probability density function

$$f_{I-H}(u; n) = \frac{1}{2(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} (u-k)^{n-1} \text{sgn}(u-k).$$

From the above, for $k \geq 7$

$$\mathbb{P}(K = k) = A_1 \begin{cases} A_2 A_3^l & 7 + 16l \leq k \leq 14 + 16l \\ A_3^l \mathbf{I}(k - 16l - 14) & 15 + 16l \leq k \leq 22 + 16l, \end{cases}$$

where l is $\lfloor (k-7)/16 \rfloor$ and A_1 , A_2 , A_3 , \mathbf{I} are given by (A.26), (A.27), (A.28), (A.29) below.

$$A_1 = \mathbb{P}(a_1 < b^{(0)}) \mathbb{P}(a_3 + a_2 < b^{(1)}) \mathbb{P}\left(\sum_{i=4}^7 a_i < b^{(2)}\right) = \frac{1}{2} \frac{1}{2} \frac{1}{2} = \frac{1}{8}. \quad (\text{A.26})$$

$$A_2 = \mathbb{P}(a_8 \geq b^{(3)}) = \mathbb{P}(a_{24} \geq b^{(4)}) = \frac{1}{16}. \quad (\text{A.27})$$

$$A_3 = \mathbb{P}\left(\sum_{i=8}^{23} a_i < b^{(3)}\right) = \int_0^8 \frac{1}{8} db^{(3)} \int_0^{b^{(3)}} f_{I-H}(u; 16) du = 0.0578. \quad (\text{A.28})$$

\mathbf{I} is an 8-element vector. The element in position $x - 14$ ($x \in [15, 22]$) is given by

$$\mathbf{I}(x-14) = \mathbb{P}\left(\sum_{i=8}^x a_i < b^{(3)} \text{ and } \sum_{i=8}^{x+1} a_i \geq b^{(3)}\right) = \int_0^1 da_{x+1} \int_0^8 \frac{1}{8} db^{(3)} \int_{b^{(3)}-a_{x+1}}^{b^{(3)}} f_{I-H}(u; x-7) du. \quad (\text{A.29})$$

We have shown the computations for $k \in [15, 23]$ and $l = 0$. For $24 \leq k \leq 38$ and $l = 1$, we have

$$\begin{aligned} \mathbb{P}(K = k) &= \mathbb{P}(a_1 < b^{(0)})\mathbb{P}(a_3 + a_2 < b^{(1)})\mathbb{P}\left(\sum_{i=4}^7 a_i < b^{(2)}\right)\mathbb{P}\left(\sum_{i=8}^{23} a_i < b^{(3)}\right) \\ &\quad \cdot \mathbb{P}\left(\sum_{i=24}^k a_i < b^{(4)} \text{ and } \sum_{i=24}^{k+1} a_i \geq b^{(4)}\right), \end{aligned}$$

due to the independence of the backoff counters. The last factor is equal to

$$\mathbb{P}\left(\sum_{i=8}^k a_i < b^{(3)} \text{ and } \sum_{i=8}^{k+1} a_i \geq b^{(3)}\right) \text{ for } 8 \leq k \leq 22,$$

thus the distribution for $23 \leq k \leq 38$ is equal to the distribution for $7 \leq k \leq 22$ multiplied by A_3 . For $39 \leq k \leq 54$ the distribution is equal to the distribution for $7 \leq k \leq 22$ multiplied by A_3^2 , and so on. This completes the proof.

B Experimental Framework

We built a testbed of 19 stations in order to compare WiFi and PLC in Chapter 2, to explore the spatio-temporal variation of PLC and to validate the capacity-estimation technique proposed in Chapter 3. In addition, we constructed a testbed in ideal channel conditions in order to validate the MAC-layer findings and models of Chapters 4 and 5.

In this appendix, we describe the experimental settings that are used to produce all measurements of this work. To provide helpful information for other researchers who would like to build similar testbeds, we also introduce guidelines for configuring and for obtaining various metrics from PLC devices.

B.1 Testbed and Setup

Our main testbed consists of Alix 2D2 boards running the Openwrt Linux distribution [60]. The boards are equipped with a HomePlug AV miniPCI card (Meconet interface, Intellon INT6300 chip) that interacts with the kernel through a Realtek Ethernet driver, and with an Atheros AR9220 wireless interface. The stations are spread over the second floor of the EPFL BC building, as shown on Figure 3.1.

As mentioned in Section 2.6, each PLC station must be authenticated by a CCo. If the station does not sense any CCo transmitting beacons, then it forms its own network. To avoid modifications in the network structure, we set the CCo statically in our testbed. We build two different PLC networks. These networks have different encryption keys (there is encryption on the MAC layer) thus, only stations belonging to the same network can communicate with each other.

In addition to using our main testbed, we experiment and validate our findings with HomePlug AV500 devices, the Netgear XAVB5101 (Atheros QCA7400 chip).

Modifying Topology and Security Key

To establish the two aforementioned PLC networks and to maintain a constant network PLC topology for our studies, we need to set statically the CCo and to setup different

Appendix B. Experimental Framework

security keys for each PLC network. To this end, we modify the platform independent binary (PIB) files of the interfaces. We use the Atheros Open Powerline Toolkit [61], a tool that interacts with the HomePlug AV chips¹. The tool uses vendor-specific management messages (MMs), as described in Section 2.5, to interact with, and to configure the devices. It also provides functions to retrieve, modify and update the PIB files. For instance, to change the CCo mode and the network membership key (NMK) of a station with MAC address <MAC>, we can execute the following commands:

```
int6k -p <PIBFILE> <MAC> # Retrieve the PIB file from the PLC interface
modpib -C <MODE> -v <PIBFILE> # Modify CCo mode of <PIBFILE>, <MODE>: 0 Auto, 1
    Never, 2 Always CCo
modpib -N <NMK> -v <PIBFILE> # Modify NMK
int6k -P <PIBFILE> <MAC> # Update the PIB file of the PLC interface
```

B.2 Measurement and Traffic Tools

To retrieve the metrics for the PHY and MAC performance evaluation, we use the Atheros Open Powerline Toolkit [61]. The tool [61] provides functions to interact with the old and the latest Atheros PLC chips, such as QCA7400².

To capture the PLC frame headers, we use `faifa` [62]. Similar to [61], `faifa` interacts with PLC devices using MMs. It enables a “sniffer” mode with which we can capture the SoF delimiters of all received PLC MPDUs.

To generate traffic, we use `iperf`. For all the experiments, links are saturated with UDP traffic (unless otherwise stated), i.e., stations transmit at maximum available rates, so that we can measure metrics such as capacity. To change the priority of the PLC frames and to conduct tests with various traffic classes, we modify the type-of-service field of the IP layer header by using `iperf`. Table B.1 outlines the metrics used throughout this work, as well as the methods used to measure them.

PLC Management Messages and Interaction with *Click*

As mentioned above, the open-source PLC tools use vendor-specific MMs to interact with the hardware. With the *Click* modular router [26], we can create elements to send and receive these MMs. Typical examples that we implement using *Click* are the capacity estimation using the MMs replacing the `int6krate` command and capturing packets by enabling the sniffer mode of PLC interfaces. We capture the MMs by using `tcpdump` to understand their structure, to infer their `MMType` (field that defines the message type), and to code the *Click elements* (simple C++ classes that process packets). The description of these MMs are also available through the source code of the tools we are using [61]. To interact with the PLC chip, MMs have to be encapsulated in an Ethernet header with *EtherType* `0x88E1`.

¹We are equipped with devices from 6 vendors. We can retrieve statistics from all devices using [61].

²For instance, `int6k` has to be replaced by `plctool` for the new QCA7400 chips.

B.3. MAC-Layer Experiments

Metric	Notation	Measured by
Arrival timestamp	t	SoF delimiter [62]
Bit-loading estimate of frame	BLE_s	SoF delimiter
Burst length	BL	SoF delimiter
Priority of the frame	$LinkID$	SoF delimiter
Source station of the frame	$STEI$	SoF delimiter
Number of remaining MPDUs for current burst	MPDUCnt	SoF delimiter
Delimiter Type: Frame, beacon, ACK, RTC/CTS, Sound	-	SoF delimiter
PB error probability	PB_{err}	<code>ampstat</code> [61]
Collided and Acknowledged frames	C_i and A_i	<code>ampstat</code>
Average BLE	BLE	<code>int6krate</code> [61]
Throughput	T	<code>iperf</code> or <code>ifstat</code>
MCS index (WiFi)	MCS	WiFi frame control

Table B.1 – Metrics and measurement methods. For more information about the SoF delimiter fields, see Section 6.4.1.1.5.2 of the IEEE 1901 standard [10].

B.3 MAC-Layer Experiments

Reproducing Ideal Channel Conditions for PLC and WiFi

The experiments of Chapters 4 and 5 are run under ideal-channel conditions to avoid hidden factors from the PHY layer, which can affect our performance evaluation. To emulate PLC ideal-channel conditions, we attach all stations to a common power-strip by using coupling devices. The miniPCI cards are powered by the bus and this enables us to conduct experiments without plugging the power-strip to electricity, thus avoiding external noise. This setting is similar to [47] and reproduces the *perfect channel* assumption. In addition to ideal-channel conditions, to avoid MMs overhead by non-transmitting stations, we only plug the number of stations required to be transmitting for the experiment.

All stations are placed in the same room, which also reproduces ideal channel conditions for our WiFi experiments. We set the wireless cards to 802.11a mode and to channel 44 to avoid interference with other wireless networks in our building.

Measuring Collision Probability for PLC

To measure collision probability, we use [61]. The command `ampstat` of [61] can reset to 0 or retrieve the number of acknowledged and collided MPDUs given the peer-station MAC address, the priority, and the flow direction (transmission or reception).

Before each experiment, we reset the statistics of the frames transmitted at all stations by using `ampstat`. At the end of the test, we request the number of collided and acknowledged frames transmitted from all the stations given the MAC address of the destination station. To obtain these statistics `ampstat` sends an MM with MMType 0xA030. We use the reply of the PLC interface. Specifically, the bytes 25-32 of this reply represent the number of acknowledged frames (included collided ones) and the bytes 33-40 represent the number of collided frames.

Appendix B. Experimental Framework

Let C_i be the number of collided frames transmitted by station i , and let A_i be the number of acknowledged frames transmitted by station i . To evaluate the collision probability in the network, we compute $\sum_{i=1}^N C_i / \sum_{i=1}^N A_i$. Observe that we divide only by the sum of the acknowledged frames not including the collided ones in this sum, as the 1901 standard allows selective acknowledgments for all the physical blocks contained in a frame [10]. When a collision occurs and the preambles of the collided frames can be decoded (due to the robust modulation with which they are transmitted), the destination acknowledges the frame, with an indication that all the physical blocks are received with errors, which yields a collision.

Short-Term Fairness Study

To study fairness, we need to know the station identifications of all transmitted PLC frames. Due to frame aggregation with PLC, we cannot employ `tcpdump` to capture the frames, as this tool only captures the Ethernet packets. The number of Ethernet packets contained in a PLC frame is variable hence, identifying the sequences and inter-arrival times of actual PLC frames is challenging. To address this challenge, we use `Faifa` [62] to capture the SoF delimiters of *PLC frames* and to retrieve useful metrics such as the ones introduced in Table B.1. `Faifa` prints all the fields of SoF, including the source TEI, the “MPDU Count”, and the arrival time of the frame, which are of interest in this experiment. The field MPDU Count is equal to the number of remaining MPDUs in the current burst transmission. MPDUs can belong to bursts and bursts contend for the medium, not individual MPDUs. Thus, we collect MPDUs with MPDU Count equal to 0, as these MPDUs are either not part of a burst or are the last MPDUs of a burst.

We also study the fairness of WiFi. We employ `tcpdump` to capture WiFi packets.

Bibliography

- [1] Consumer Interest in Wired Solutions to Wireless Home Networking Problems: Quantitative Findings. <http://www.mocalliance.org/news/Parks-Associates-Final-Report.pdf>. [Online; accessed June 9, 2016].
- [2] Info-graphic by Goldman Sachs: What is the Internet of Things? <http://www.goldmansachs.com/our-thinking/pages/iot-infographic.html>. [Online; accessed June 9, 2016].
- [3] Cisco: How Many Internet Connections are in the World? Right. Now. <http://blogs.cisco.com/news/cisco-connections-counter>. [Online; accessed June 9, 2016].
- [4] IEEE Standard for a Convergent Digital Home Network for Heterogeneous Technologies. *IEEE Std 1905.1-2013*, 2013.
- [5] nVoy Hybrid Networking. <http://www.nvoy.org/>. [Online; accessed June 9, 2016].
- [6] Value Proposition (Definition). <http://www.investopedia.com/terms/v/valueproposition.asp?layout=orig>. [Online; accessed June 9, 2016].
- [7] Multimedia over Coax Alliance (MoCA). <http://www.mocalliance.org/>. [Online; accessed June 9, 2016].
- [8] HomePlug Alliance. <http://www.homeplug.org/>. [Online; accessed June 9, 2016].
- [9] HomePlug Technology Overview. <http://www.homeplug.org/tech-resources/techoverview/>. [Online; accessed June 9, 2016].
- [10] IEEE Standard for Broadband over Power Line Networks: Medium Access Control and Physical Layer Specifications. *IEEE Std 1901-2010*, 2010.
- [11] IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications. *IEEE Std 1901.2-2013*, 2013.
- [12] Ramanujan K Sheshadri and Dimitrios Koutsonikolas. Comparison of routing metrics in 802.11n wireless mesh networks. In *Proc. of IEEE INFOCOM*, pages 1869–1877, 2013.

Bibliography

- [13] Richard Draves, Jitendra Padhye, and Brian Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *Proc. of ACM MobiCom*, pages 114–128, 2004.
- [14] Saumitra M Das, Himabindu Pucha, Konstantina Papagiannaki, and Y Charlie Hu. Studying wireless routing link metric dynamics. In *Proc. of ACM IMC*, pages 327–332, 2007.
- [15] Giuseppe Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, 2000.
- [16] Venkatesh Ramaiyan, Anurag Kumar, and Eitan Altman. Fixed point analysis of single cell IEEE 802.11e WLANs: uniqueness and multistability. *IEEE/ACM Transactions on Networking*, 16(5):1080–1093, 2008.
- [17] Sabih Guzelgöz, Hasan B Çelebi, Tayyar Güzel, Hüseyin Arslan, and M Kivanc Mihçak. Time frequency analysis of noise generated by electrical loads in PLC. In *Proc. of IEEE ICT*, pages 864–871, 2010.
- [18] Haniph A. Latchman, Srinivas Katar, Lawrence W. Yonge, and Sherman Gavette. *Homeplug AV and IEEE 1901*. John Wiley and Sons, Inc., 2013.
- [19] José Antonio Cortés, Luis Díez, Francisco Javier Cañete, and Jesús López. Analysis of the periodic impulsive noise asynchronous with the mains in indoor PLC channels. In *Proc. of IEEE ISPLC*, pages 26–30, 2009.
- [20] Salvador Sancha, Francisco Canete, Luis Diez, and Jose T. Entrambasaguas. A channel simulator for indoor power-line communications. In *Proc. of IEEE ISPLC*, pages 104–109, 2007.
- [21] IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks-Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *IEEE Std 802.11-1997*, 1997.
- [22] Andreas Schwager. An Overview of the HomePlug AV2 Technology. *Journal of Electrical and Computer Engineering*, 2013.
- [23] Srinivas Katar, B Mashbum, Kaywan Afkhamie, Haniph Latchman, and R Newrnan. Channel adaptation based on cyclo-stationary noise characteristics in PLC systems. In *Proc. of IEEE ISPLC*, pages 16–21, 2006.
- [24] Rohan Murty, Jitendra Padhye, Ranveer Chandra, A Roy Chowdhury, and Matt Welsh. Characterizing the end-to-end performance of indoor powerline networks. Technical report, Harvard University Microsoft Research Technical Report, 2008.
- [25] Christian Benvenuti. *Understanding Linux network internals*. O’Reilly Media, Inc., 2006.

-
- [26] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M Frans Kaashoek. The Click modular router. *ACM TOCS*, 18(3):263–297, 2000.
- [27] IEEE Standard for Information technology - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput. *IEEE Std 802.11n-2009*, 2009.
- [28] Nico Weling. Field analysis of 40.000 plc channels to evaluate the potentials for adaptive transmit power management. In *Proc. of IEEE ISPLC*, pages 201–206, 2010.
- [29] Zhang Liu, Alaeddine El Fawal, and J-Y Le Boudec. Coexistence of Multiple HomePlug AV Logical Networks: A Measurement Based Study. In *Proc. of IEEE GLOBECOM*, pages 1–5, 2011.
- [30] Anurag Kumar, Eitan Altman, Daniele Miorandi, and Munish Goyal. New Insights From a Fixed-Point Analysis of Single Cell IEEE 802.11 WLANs. *IEEE/ACM Transactions on Networking*, 15(3):588–601, 2007.
- [31] Brad Zarikoff and David Malone. Construction of a PLC testbed for network and transport layer experiments. In *Proc. of IEEE ISPLC*, pages 135–140, 2011.
- [32] Yu-Ju Lin, Haniph A Latchman, Richard E Newman, and Srinivas Katar. A comparative performance study of wireless and power line networks. *IEEE Communications Magazine*, 41(4):54–63, 2003.
- [33] Peerapol Tinnakornsriruphap, Punyaslok Purkayastha, and Bibhu Mohanty. Coverage and capacity analysis of hybrid home networks. In *Proc. of IEEE ICNC*, pages 117–123, 2014.
- [34] Paul P. Patras, Albert Banchs, Pablo Serrano, and Arturo Azcorra. A control-theoretic approach to distributed optimal configuration of 802.11 WLANs. *IEEE Transactions on Mobile Computing*, 10(6):897–910, 2011.
- [35] Giuseppe Bianchi, Luigi Fratta, and Matteo Oliveri. Performance evaluation and enhancement of the CSMA/CA MAC protocol for 802.11 wireless LANs. In *Proc. of IEEE PIMRC*, pages 392–396, 1996.
- [36] Min Young Chung, Myoung-Hee Jung, Tae-Jin Lee, and Yutae Lee. Performance analysis of HomePlug 1.0 MAC with CSMA/CA. *IEEE Journal on Selected Areas in Communications*, 24(7):1411–1420, 2006.
- [37] Open source FirmWare for WiFi networks (OpenFWWF). <http://netweb.ing.unibs.it/~openfwf/>. [Online; accessed June 9, 2016].

Bibliography

- [38] Pablo Salvador, Vincenzo Mancuso, Pablo Serrano, Francesco Gringoli, and Albert Banchs. VoIPiggy: Analysis and Implementation of a Mechanism to Boost Capacity in IEEE 802.11 WLANs Carrying VoIP traffic. *IEEE Transactions on Mobile Computing*, 13(7):1640–1652, 2014.
- [39] Pablo Salvador, Luca Cominardi, Francesco Gringoli, and Pablo Serrano. A First Implementation and Evaluation of the IEEE 802.11aa Group Addressed Transmission Service. *ACM SIGCOMM Comput. Commun. Rev.*, 42(1):13–18, 2014.
- [40] The Multi-Generator Toolset. <http://cs.itd.nrl.navy.mil/work/mgen/>. [Online; accessed June 9, 2016].
- [41] Cristina Cano and David Malone. On Efficiency and Validity of Previous Homeplug MAC Performance Analysis. *Computer Networks*, 83(C):118–135, 2015.
- [42] Giuseppe Bianchi and Ilenia Tinnirello. Kalman filter estimation of the number of competing terminals in an IEEE 802.11 network. In *Proc. of IEEE INFOCOM*, pages 844–852, 2003.
- [43] Miguel Elias M Campista, LHMK Costa, and OCMB Duarte. Improving the multiple access method of CSMA/CA home networks. In *Proc. of IEEE CCNC*, pages 645–649, 2006.
- [44] Evan Kriminger and Haniph Latchman. Markov chain model of homeplug CSMA MAC for determining optimal fixed contention window size. In *Proc. of IEEE ISPLC*, pages 399–404, 2011.
- [45] Cristina Cano and David Malone. When priority resolution goes way too far: an experimental evaluation in PLC networks. In *Proc. of IEEE ICC*, pages 952–957, 2015.
- [46] Cristina Cano and David Malone. Performance evaluation of the priority resolution scheme in PLC networks. In *Proc. of IEEE ISPLC*, pages 290–295, 2014.
- [47] Brad Zarikoff and David Malone. Construction of a PLC test bed for network and transport layer experiments. In *Proc. of IEEE ISPLC*, pages 135–140, 2011.
- [48] Brad Zarikoff and David Malone. Dynamics of contention free period reservation in IEEE 1901 networks. In *Proc. of IEEE ISPLC*, pages 242–247, 2012.
- [49] Gilles Berger-Sabbatel, Andrzej Duda, Olivier Gaudoin, Martin Heusse, and Franck Rousseau. Fairness and its impact on delay in 802.11 networks. In *Proc. of IEEE GLOBECOM*, pages 2967–2973, 2004.
- [50] Michael Bredel and Markus Fidler. Understanding Fairness and its Impact on Quality of Service in IEEE 802.11. In *Proc. of IEEE INFOCOM*, pages 1098–1106, 2009.

- [51] Raj Jain, Dah-Ming Chiu, and William R Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. *DEC Research Report TR-301*, 1984.
- [52] Can Emre Koksal, Hisham Kassab, and Hari Balakrishnan. An analysis of short-term fairness in wireless media access protocols (poster session). In *Proc. of ACM SIGMETRICS*, pages 118–119, 2000.
- [53] Stuart A Klugman, Harry H Panjer, and Gordon E Willmot. *Loss models: from data to decisions*, volume 2. Wiley New York, 1998.
- [54] Jeong-woo Cho, J-Y Le Boudec, and Yuming Jiang. On the asymptotic validity of the decoupling assumption for analyzing 802.11 MAC protocol. *IEEE Transactions on Information Theory*, 58(11):6879–6893, 2012.
- [55] Kaidi Huang, Ken R Duffy, and David Malone. On the Validity of IEEE 802.11 MAC Modeling Hypotheses. *IEEE/ACM Transactions on Networking*, 18(6):1935–1948, 2010.
- [56] Charles Bordenave, David McDonald, and Alexandre Proutiere. A particle system in interaction with a rapidly varying environment: mean field limits and applications. *Networks and Heterogeneous Media*, 5(1):31–62, 2010.
- [57] Gaurav Sharma, Ayalvadi Ganesh, and Peter Key. Performance Analysis of Contention Based Medium Access Control Protocols. *IEEE Transactions on Information Theory*, 55(4):1665 –1682, 2009.
- [58] Michel Benaïm and Jean-Yves Le Boudec. A class of mean field interaction models for computer and communication systems. *Performance Evaluation*, 65(11–12):823 – 838, 2008.
- [59] Ken R Duffy. Mean field Markov models of wireless local area networks. *Markov Processes and Related Fields*, 16(2):295–328, 2010.
- [60] OpenWrt. <https://openwrt.org/>. [Online; accessed June 9, 2016].
- [61] Qualcomm Atheros. Open Powerline Toolkit. <https://github.com/qca/open-plc-utils>. [Online; accessed June 9, 2016].
- [62] Faifa. <http://github.com/ffainelli/faifa>. [Online; accessed June 9, 2016].

Christina Vlachou

✉ christina.g.vlachou@gmail.com

🏠 people.epfl.ch/vlachou

Education

Sep. 2011 – **PhD**, *École Polytechnique Fédérale de Lausanne (EPFL)*, Switzerland,
Jun. 2016 Advisor: Patrick Thiran.

Thesis: Measuring, Modeling and Enhancing Power-Line Communications (PLC)

Sep. 2006 – **Diploma in Electrical and Computer Engineering**, *National Technical University of Athens (NTUA)*, Greece, Advisor: P. G. Cottis.

Jul. 2011

GPA: 9.22/10.00, Class rank 11 out of 325

Thesis: Network Coding for Overhead Reduction in Delay Tolerant Networks (DTNs)

- Innovative algorithm for overhead reduction and buffer management under epidemic routing protocols in DTNs

Experience

Sep. 2011 – **Research Assistant**, *Laboratory for Communications and Applications at EPFL*,
Jul. 2016 Lausanne, Switzerland.

Research on the dominant PLC standard – which has not received any attention by the research community, although it is widely used – with emphasis on the MAC layer:

- Introduced the first analytical models of the PLC CSMA/CA protocol
 - Enhanced PLC throughput by 27% compared to commercial devices
 - Unveiled that the PLC CSMA/CA is short-term unfair, which yields high jitter
 - Uncovered that the CSMA/CA parameters tune a tradeoff between throughput and jitter
 - Designed an algorithm that minimizes jitter for delay-sensitive applications
 - Pioneered PLC experimental studies and validated all findings via measurements
- Design and testbed implementation of hybrid Wi-Fi/PLC networks:
- Thoroughly investigated spatio-temporal variations of performance
 - Introduced systematic guidelines for accurate link-metric estimation
 - Routing/congestion-control algorithms yield up to 10x gains compared to using only Wi-Fi

Jun. 2015 – **Research Intern**, *Qualcomm R&D*, San Diego, USA.

- Sep. 2015
- Rate prediction by using channel quality feedback in dense Wi-Fi networks
 - Interference analysis in dense Wi-Fi networks (IEEE 802.11ax)
 - Traffic models implementation on C++ simulator

Sept. 2013 – **Firmware-Design Intern**, *Marvell*, Etoy, Switzerland.

- Oct. 2013
- Firmware design for PLC: implementation of functions for PHY layer management
 - Extensive FPGA-platform experiments for firmware debugging and hardware interoperability

Awards and Honors

- 2015 **Best Use of Vuforia**, “HackMobile” intern hackathon, Qualcomm.
Developed (within 16 hours) an application for augmented-reality 3D visualization of music by using Vuforia platform, without having prior background on the technologies used
- 2014 **Best Paper Runner-up**, *IEEE International Conference on Network Protocols (ICNP)*.
- 2011 **1st Year Fellowship**, *Computer and Communication Sciences Doctoral School, EPFL*.
- 2007 **Papakyriakopoulos Award for Excellence in Mathematics**, NTUA.
Awarded to the second-year student with the highest grades in Mathematics courses

Publications

1. **Christina Vlachou**, Albert Banchs, Julien Herzen, and Patrick Thiran. How CSMA/CA with Deferral Affects Performance and Dynamics in Power-Line Communications. *IEEE/ACM Transactions on Networking (to appear)*, 2017.
2. **Christina Vlachou**, Albert Banchs, Pablo Salvador, Julien Herzen, and Patrick Thiran. Analysis and Enhancement of CSMA/CA with Deferral in Power-Line Communications. *IEEE Journal of Selected Areas in Communications*, 2016.
3. **Christina Vlachou**, Sébastien Henri, and Patrick Thiran. Electri-Fi Your Data: Measuring and Combining Power-Line Communications with Wi-Fi. In *Proc. ACM Internet Measurement Conf. (IMC)*, 2015.
4. **Christina Vlachou**, Albert Banchs, Julien Herzen, and Patrick Thiran. Analyzing and Boosting the Performance of Power-Line Communication Networks. In *Proc. ACM International Conf. on emerging Networking EXperiments and Technologies (CoNEXT)*, 2014.
5. **Christina Vlachou**, Albert Banchs, Julien Herzen, and Patrick Thiran. On the MAC for Power-Line Communications: Modeling Assumptions and Performance Tradeoffs. In *Proc. IEEE International Conf. on Network Protocols (ICNP)*, 2014.
6. **Christina Vlachou**, Albert Banchs, Julien Herzen, and Patrick Thiran. Performance Analysis of MAC for Power-Line Communications. In *Proc. ACM SIGMETRICS Conf. (poster session)*, 2014.
7. **Christina Vlachou**, Julien Herzen, and Patrick Thiran. Fairness of MAC protocols: IEEE 1901 vs. 802.11. In *Proc. IEEE International Symp. Power Line Communications and Its Applications (ISPLC)*, 2013.
8. Anna Vazintari, **Christina Vlachou**, and Panayotis G Cottis. Network Coding for Overhead Reduction in Delay Tolerant Networks. *Wireless personal communications*, 72(4), Springer 2013.

Skills

Network protocols	Extensive knowledge of the TCP/IP stack, Analytical modeling and performance evaluation of network protocols, Hands-on experience with network testbeds.	Platforms for networks	Click Modular Router, ns-3, Wireshark, Atheros PLC toolkit
Coding	C++, Matlab, C, Python, Bash scripting, Qt	Standards	Extensive knowledge of IEEE 1901, 1905, 802.11

Teaching Experience

- Fall 2015 **Research Tutorial**, *Bachelor's class: Computer Networks*, EPFL.
Tutorial on the CSMA/CA protocol of PLC.
- Falls '12-'15 **Teaching Assistant**, *Dynamical Systems Theory for Engineers*, EPFL.
- Spring 2014 **Teaching Assistant**, *Probability and Statistics*, EPFL.
- Spring 2013 **Teaching Assistant**, *Introduction to Programming*, EPFL.
- Fall 2015 **Project Supervision**, *Suresh Aswin: Master's project*, EPFL.
Simulation of hybrid Wi-Fi/PLC networks
- Spring 2015 **Project Supervision**, *Ting-Hsuan Lee: Master's project*, EPFL.
Capacity estimation and load balancing in hybrid networks
- Summer 2013 **Project Supervision**, *Yahya H. Ezzeldin: Intern at Summer@EPFL program*, EPFL.
Evaluation of PLC and Wi-Fi links for hybrid in-home multi-hop transmission

Professional Services

- External Reviewer** *SIGMETRICS 2016, NSDI 2015, CoNEXT 2013*
- Reviewer** *IEEE JSAC, IEEE Communications Letters, IEEE/ACM Transactions on Networking*
- TPC Member** *IEEE GLOBECOM 2016 - Selected Areas in Communications: Power Line Communications*

Extracurricular Activities

- 2015–2016 **President of AEGEL**, *EPFL's official association of Greek students in Lausanne.*
- Summer 2015 **Top 10 in IdeaQuest**, *Qualcomm's intern innovation challenge: built and demoed a smart chair that corrects the user posture by using Arduinos and pressure sensors.*