Space Networking: Models and Methods

Juan Andrés A Fraire

▶ **To cite this version:**

Juan Andrés A Fraire. Space Networking: Models and Methods. Computer Science [cs]. INSA Lyon, 2025. ⟨tel-05059552⟩

HAL Id: tel-05059552

https://hal.science/tel-05059552v1

Submitted on 7 May 2025

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

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

HAL Authorization

INSA LYON

HABILITATION À DIRIGER DES RECHERCHES

Space Networking

Models and Methods

Juan A. Fraire

Jury:

M. POSS Michël, Université de Montpellier, Rapporteur
M. POPOVSKI Petar, Aalborg University, Rapporteur
MME. UYSAL Elif, Middle East Technical University, Rapporteur
MME. KARABULUT KURT Gunes, Polytechnique Montréal, Examinatrice
MME. RATNASAMY Sylvia, University of California, Berkeley, Examinatrice
M. SHENKER Scott, University of California, Berkeley, Examinateur
MME. GUERIN-LASSOUS Isabelle, Université Claude Bernard Lyon 1, Examinatrice
M. RIVANO Hervé, INSA Lyon, Examinateur

May 7, 2025

To Paula, who grounds my orbit,
and to Benicio, already fluent
in the language of the stars.

Foreword

Preparing this Habilitation to Conduct Research (HDR) has been a profoundly reflective and rewarding journey. It required significant time and energy, but it was undeniably worth it. I strongly encourage my colleagues to undertake a similar experience if given the opportunity.

One of the most fulfilling aspects of this process has been revisiting, selecting, and categorizing my research contributions. This intellectually stimulating act of organization, uniquely human despite AI's growing capabilities, led to reviving inspiring discussions with my supervisors and students. These moments provided fresh perspectives on past work, revealed the evolution of my thinking over time, and helped me to clarify my role and forward-looking vision.

Writing this HDR felt akin to constructing a language—an organized system of thought and meaning. This resonates with my admiration for French culture, particularly its profound relationship with language. As Jacques Lacan emphasized, language mediates our relationship with reality, structuring our understanding of the world. This notion is particularly significant in space networks, where intricated inter-dependent phenomena lie beyond direct observation and must be interpreted through limited telemetry or simulations. By harnessing cutting-edge yet applicable computer science models and methodologies, this HDR signifies my contribution to the scientific community, offering a unique framework for understanding and navigating the unseen realm of space networking.

It brings me immense joy to think colleagues or students might find value in these reflections. I hope this work, in its small way, advances our collective understanding of a "language" for space networking, paving the way for exploring uncharted spacetime domains.

With gratitude,
Juan A. Fraire

Abstract

This document synthesizes my contributions to space networking over the last decade, offering a cohesive framework for addressing fundamental challenges in this field.

It encompasses space network applications, architectures, environment, analysis, design, operations, routing, flow control, links, and medium access control, all tackled using advanced computer science *models* and *methods* to address the unique spatial and temporal dynamics of space and satellite systems. Notable contributions in my career include advancements in delay-tolerant networking, contact plan design, routing algorithms tailored to dynamic and constrained environments, and the integration of IoT with satellite networks. Despite the significant challenges, many of these contributions have been designed to focus on applicability. Many of them have been tested in orbit, including routing code currently running on the ISS, autonomous operations without human intervention on the GomX-4 mission, and the validation of store-and-forward protocols on ESA's OPS-SAT mission, among others.

These achievements would not have been possible without the unwavering support of inspiring mentors, the dedication of talented students, and productive collaborations. Together, they have fostered the emergence of an informal yet impactful cross-border and cross-institutional team devoted to advancing space informatics and networking. While this foundation sets the stage for an exciting next chapter of my career, it also introduces significant challenges, which I address at the end of this document.

By bridging theoretical models with practical in-orbit implementations, this HDR lays a foundation for scalable, robust, and efficient communication systems that extend from Earth to deep space and beyond.

Outline

The Habilitation to Conduct Research (HDR) is the highest academic qualification in France, signifying a researcher's ability to independently lead and supervise doctoral research.

This manuscript is submitted as part of the HDR qualification process. It synthesizes years of impactful research, showcasing originality, innovation, and leadership in the candidate's field, while outlining future research perspectives.

Chapters This document consists of ≈**60 pages of core content**, structured in three chapters to succinctly present my trajectory, principal contributions, and future perspectives. Figure 1 illustrates the organization of this document. The core chapters include the essential content required for the HDR evaluation. These are as follows:

- Chapter 1: **Background and Trajectory** – An overview of my career path, publications, funding, and research network.
- Chapter 2: **Models and Methods** – A summary of how my research has contributed to the evolution of the space networking field.
- Chapter 3: **Perspectives and Vision** – An outline of short- and long-term space networking research directions and challenges.

Appendices Chapters are followed by ≈**140 pages of optional appendices**, offering supplementary research explanations for on-demand consultation. The supporting appendices provide deeper insights into specific aspects of my research summarized in Chapter 2. These are as follows:

- Appendix A: **Applications, Architectures, and Environment** – Addressing foundational questions about the utilization and characteristics of space systems.
- Appendix B: **Analysis, Design, and Operations** – Proposing methodologies for analyzing, designing, and managing space networks.

- Appendix C: **Routing and Flow Control** – Contributing to multi-hop routing strategies and data flow management in satellite networks.
- Appendix D: **Links and Medium Access Control** – Exploring efficient single-hop data transfer and coordination mechanisms in space communication systems.

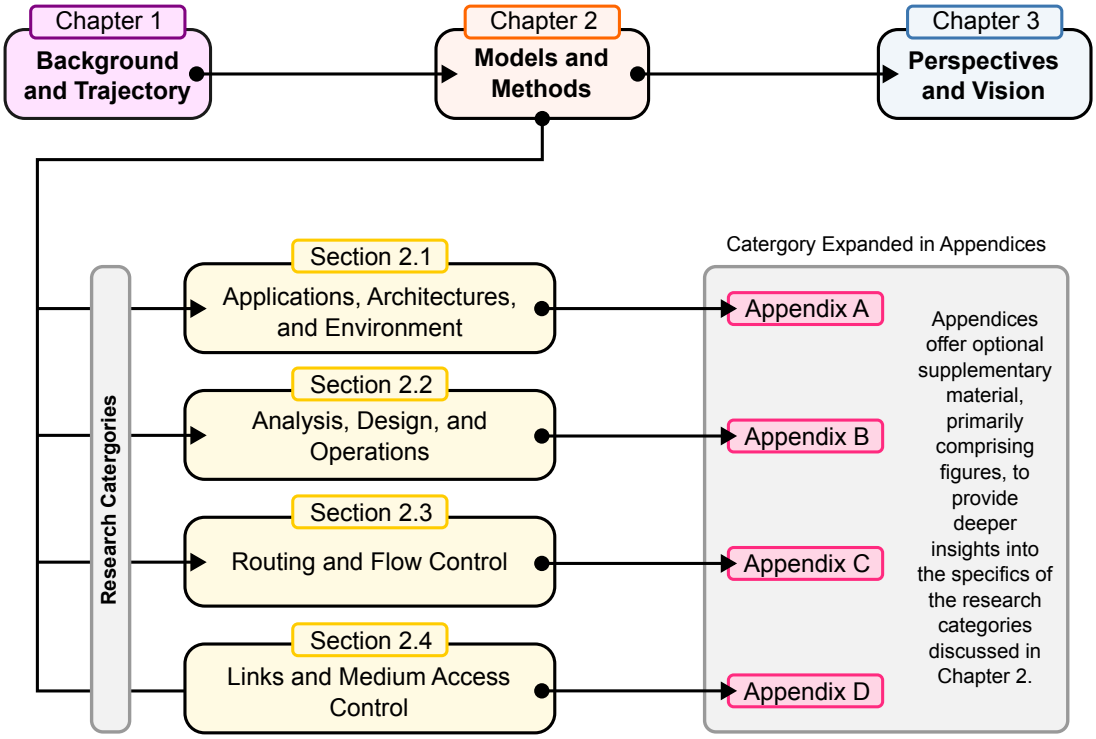

Figure 1: HDR Outline.

Reading Although extensive, nearly half of this document consists of figures (particularly in the appendices) that **can be skipped without interrupting the narrative flow**. The textual content is designed for readability, and generous fonts, margins, and spacing are used thoughtfully to ensure a comfortable reading experience.

Bibliography This HDR cites 120 co-authored papers over ≈**30 pages of bibliography**. To maintain focus and avoid an unrealistically long document, **I have not included additional citations for related works**. Instead, the reader is addressed to the most relevant references in the cited papers. These references allow direct access to foundational previous studies through convenient DOI and HAL links.

Contents

Foreword 1

Abstract 2

Outline 3

1 Background and Trajectory 10
- 1.1 Background . . . 10
- 1.2 Motivation . . . 11
- 1.3 Trajectory . . . 12
- 1.4 Publications . . . 14
  - 1.4.1 Categories . . . 15
  - 1.4.2 Methodologies . . . 17
  - 1.4.3 Roles . . . 17
  - 1.4.4 Venues . . . 18
- 1.5 Funding . . . 20
  - 1.5.1 Active Funding . . . 20
  - 1.5.2 Concluded Funding . . . 23
- 1.6 Group . . . 25
- 1.7 Reflection . . . 28

2 Models and Methods 30
- 2.1 Applications, Architectures, and Environment . . . 30
  - 2.1.1 Applications . . . 32
    - 2.1.1.1 Space Network Services . . . 32
    - 2.1.1.2 Space Cloud Services . . . 33
  - 2.1.2 Architectures . . . 34
    - 2.1.2.1 Near-Earth Networks . . . 34
    - 2.1.2.2 Deep-Space Networks . . . 35
  - 2.1.3 Environment . . . 37
    - 2.1.3.1 Radiation Exposure Campaigns . . . 37

2.1.3.2 Fault Injection Campaigns . . . 38
2.1.3.3 Applications in Networks on Chip . . . 38
2.2 Analysis, Design, and Operations . . . 39
2.2.1 Constellation Analysis and Design . . . 39
2.2.1.1 Constellation Analysis . . . 39
2.2.1.2 Constellation Design . . . 40
2.2.2 Contact Plan Design . . . 41
2.2.2.1 CPD Methodologies . . . 41
2.2.2.2 Tooling: Contact Plan Designer . . . 42
2.2.3 Mission Operation . . . 43
2.2.3.1 Automated Battery-Aware Operations . . . 43
2.2.3.2 Advanced Operation Models . . . 43
2.3 Routing and Flow Control . . . 45
2.3.1 Low-Earth Orbit Constellations . . . 45
2.3.1.1 Routing . . . 45
2.3.1.2 Flow Control . . . 46
2.3.2 Delay-Tolerant Networks . . . 46
2.3.2.1 Routing with Scheduled Contact Plans . . . 47
2.3.2.2 Routing with Uncertain Contact Plans . . . 49
2.3.2.3 Tooling: DtnSim . . . 50
2.3.2.4 Flow Control . . . 50
2.4 Links and Medium Access Control . . . 52
2.4.1 Links . . . 52
2.4.1.1 Radio Frequency (RF) Links . . . 52
2.4.1.2 Free-Space Optical (FSO) Links . . . 53
2.4.2 Medium Access Control . . . 54
2.4.2.1 Direct-to-Satellite IoT Access . . . 54
2.4.2.2 Direct-to-Cell Access . . . 56

3 Perspectives and Vision 58
3.1 Perspectives . . . 58
3.1.1 On Applications, Architectures, and Environment . . . 58
3.1.2 On Analysis, Design, and Operations . . . 59
3.1.3 On Routing and Flow Control . . . 61
3.1.4 On Links and Medium Access Control . . . 62
3.1.5 Future Projects . . . 63
3.2 Vision . . . 64
3.2.1 Scientific Vision . . . 64
3.2.2 Applied Vision . . . 66

3.2.3 Methodological Vision . . . 68
3.2.4 Vision Risks . . . 69

Appendices **71**

A Applications, Architectures, and Environment **72**
A.1 Applications . . . 72
A.1.1 Space Network Services . . . 73
A.1.1.1 Observation Data Transport . . . 73
A.1.1.2 User Data Transport . . . 74
A.1.2 Space Cloud Services . . . 79
A.1.2.1 Server Orchestration . . . 80
A.1.2.2 Server Placement . . . 80
A.1.2.3 Tooling: MeteorNet . . . 81
A.2 Architectures . . . 83
A.2.1 Near-Earth Networks . . . 83
A.2.1.1 High-Altitude Platforms . . . 83
A.2.1.2 Direct-to-Satellite IoT . . . 84
A.2.1.3 Tooling: FLoRaSat . . . 86
A.2.2 Deep-Space Networks . . . 88
A.2.2.1 Interplanetary . . . 89
A.2.2.2 Tooling: IPN-V . . . 91
A.2.2.3 Stacks: ION and μD3TN . . . 94
A.2.2.4 Interstellar . . . 97
A.3 Environment . . . 97
A.3.1 Radiation Exposure Campaigns . . . 99
A.3.1.1 Methodology . . . 100
A.3.1.2 Memory Characterization . . . 100
A.3.2 Fault Injection Campaigns . . . 101
A.3.2.1 Methodology . . . 101
A.3.2.2 Processors Characterization . . . 102
A.3.3 Networks on Chip Applications . . . 102

B Analysis, Design, and Operations **105**
B.1 Constellation Analysis and Design . . . 105
B.1.1 Constellation Analysis . . . 105
B.1.1.1 Dense Constellations . . . 106
B.1.1.2 Sparse Constellations . . . 107
B.1.2 Constellation Design . . . 111

B.1.2.1 Single-Satellite Missions . . . 111
B.1.2.2 Multi-Satellite Missions . . . 113
B.2 Contact Plan Design . . . 115
B.2.1 Problem Definition . . . 118
B.2.2 Contact Plan Design Methods . . . 120
B.2.2.1 Topology-Awareness . . . 120
B.2.2.2 Route-Awareness . . . 121
B.2.2.3 Traffic-Awareness . . . 123
B.2.2.4 Battery-Awareness . . . 124
B.2.2.5 Position-Awareness . . . 129
B.2.2.6 Tooling: Contact Plan Designer . . . 134
B.3 Mission Operation . . . 135
B.3.1 Automated Battery-Aware Operations . . . 135
B.3.2 Advanced Operation Models . . . 139
B.3.2.1 Polygon-Based Coverage Models . . . 139
B.3.2.2 Network Size Estimators . . . 139
B.3.2.3 Satellite Visibility Predictors . . . 143
B.3.2.4 Parallelizing Orbital Propagators . . . 143

C Routing and Flow Control 145
C.1 Low-Earth Orbit Constellations . . . 145
C.1.1 Routing . . . 146
C.1.1.1 OSPF for Satellite Networks . . . 146
C.1.1.2 DiscoRoute . . . 148
C.1.2 Flow Control . . . 148
C.1.2.1 Evaluation of BBR in QUIC . . . 150
C.1.2.2 Congestion Control and Route Stability . . . 150
C.2 Delay-Tolerant Networks . . . 151
C.2.1 Routing with Scheduled Contact Plans . . . 153
C.2.1.1 Contact Graph Routing . . . 154
C.2.1.2 Reliability . . . 158
C.2.1.3 Applicability . . . 158
C.2.1.4 Viability . . . 159
C.2.1.5 Scalability . . . 161
C.2.1.6 Shortest-Path Tree Routing . . . 165
C.2.2 Routing with Uncertain Contact Plans . . . 167
C.2.2.1 CGR-2Copies . . . 168
C.2.2.2 Markov Decision Processes . . . 168
C.2.2.3 Statistical Approaches . . . 171

C.2.2.4 Learning Approaches . . . 173
C.2.3 Tooling: DtnSim . . . 173
C.2.4 Flow Control . . . 175
C.2.4.1 Reliable Delivery . . . 177
C.2.4.2 Congestion Control . . . 178

D Links and Medium Access Control 185
D.1 Links . . . 185
D.1.1 Radio Frequency (RF) Links . . . 186
D.1.1.1 Proximity-1 Modem . . . 186
D.1.1.2 Contact Determination and Tracing . . . 188
D.1.2 Free-Space Optical (FSO) Links . . . 189
D.1.2.1 Space-to-Ground . . . 189
D.1.2.2 Space-to-Space . . . 190
D.2 Medium Access Control . . . 194
D.2.1 Direct-to-Satellite IoT Access . . . 195
D.2.1.1 Analysis . . . 195
D.2.1.2 Enhancement . . . 196
D.2.1.3 Transmission Policies . . . 200
D.2.1.4 Custom Protocol Design . . . 203
D.2.1.5 Management . . . 206
D.2.1.6 DTN Integration . . . 208
D.2.2 Direct-to-Cell Access . . . 210

Chapter 1

Background and Trajectory

This chapter overviews my research **Background** and the **Motivation** driving its investigation at this pivotal moment. It further details my career **Trajectory**, key **Publications**, secured **Funding**, and the collaborative **Group** behind the research presented throughout this document.

1.1 Background

Humanity has always gazed into the skies with an innate desire to explore. The act of looking beyond, of imagining what lies above and beyond the stars, seems intrinsic to our nature. Not long after launching the first satellite into orbit, we ventured to the Moon and, soon after, even further into the cosmos. Landing robots on other planets and exploring the mysteries of our universe remain among the deepest fantasies shared by dreamers, scientists, and explorers alike.

> Realizing these ambitions requires mastery of space technology's intricate and cross-disciplinary science. To justify the immense resources demanded by these endeavors, we must also **uncover applications that provide tangible value** to our terrestrial lives—a synergy between exploration and service that strengthens humanity's resolve to push the boundaries of the unknown.

Today, the space sector is thriving. We stand on the cusp of a new era reminiscent of the 1990s when large-scale constellations first emerged but faltered due to unsustainable business models. However, this second wave is poised to bring a distinct transformation. The infusion of private capital, geopolitical competition, the pervasive need for global internet access, and transformative informatics and computer science advancements drive its momentum. Technologies like artificial intelligence, mere dreams in the 1990s, now reshape what is possible, making this moment uniquely opportune. This

confluence of innovation, ambition, and necessity marks **a perfect time to craft new solutions for space networking**, uniting the aspirations of exploration with the practical benefits they bring to modern applications and everyday terrestrial life.

As detailed in this HDR, my work finds its purpose in this dynamic landscape, poised to contribute to this exciting space exploration and innovation chapter.

1.2 Motivation

Informatics in space is a niche discipline, quietly critical yet historically underserved. However, it rapidly evolves, with advancements scaling toward a more significant and impactful future. Artificial intelligence, in its broadest sense, is thriving, reshaping disciplines across science and technology—and space informatics is no exception. The convergence of these fields creates unparalleled opportunity, making this an ideal time to formalize and contribute to this burgeoning area of research. Succeeding in this endeavor will pave the way for effective and efficient exploration of the Solar System and beyond.

> As near-Earth networked satellite systems scale to thousands of nodes, analyzing, designing, and operating them surpasses human capabilities. In this context, **integrating applicable and validated computer models and methods becomes indispensable**, with the rise of AI arriving at a pivotal moment. However, AI in this context must be interpreted broadly, incorporating **precise, explainable, and deterministic techniques**. The margin for randomness is virtually nonexistent in complex and costly space missions.

This work is not merely aligned with current trends; it aims to shape and define them. Concepts I introduced to the community, such as *contact plan design* and *direct-to-satellite IoT*, which were once abstract ideas, are now gaining traction in the scientific and engineering community. Developing such foundational frameworks is no simple task, but it underscores the importance of establishing a precise vocabulary and methodology for navigating the complexities of space networking. Most importantly, **grounding models and methods in realistic assumptions and validating them in orbit whenever possible are crucial steps** in transferring these advancements to space companies and agencies, enabling their practical realization.

This HDR represents a pivotal step in consolidating my contributions to this transformative field. It advances its theoretical underpinnings while strongly focusing on practical applications with concrete value.

1.3 Trajectory

I am an Argentinian researcher who specializes in space and satellite communication networks. My professional trajectory, depicted in Figure 1.1, has included positions at several prestigious institutions in Europe and South America.

Studies I earned my Telecommunications Engineering degree from the Instituto Universitario Aeronáutico (IUA), Argentina, in 2008. Recognized for my academic excellence through multiple awards, Motorola promptly recruited me as a New Product Engineer. There, I gained hands-on experience with emerging wireless mobile networks and **valuable insights into the industry's practical and strategic costumes**. This experience revealed my true calling: addressing the most fundamental questions in technology.

Research Following this, I moved to the academic sector to pursue my PhD at Universidad Nacional de Córdoba (UNC), supervised by Prof. Pablo A. Ferreyra and Prof. Jorge M. Finochietto. I graduated in 2015 with my dissertation focused on **applying computer science methods to resource-constrained space networks**. This research was framed in the SARE mission from CONAE (the Argentinian Space Agency). This ensured that my work addressed concrete and practical challenges in real-world scenarios.

In 2016, I commenced my postdoctoral research at the TIMA Laboratory from the Université Grenoble Alpes (UGA) in France, where I investigated advanced **reliability mechanisms for electronic space systems**. This experience, supervised by Dr. Raoul Velazco, immersed me in the European research culture and deepened my understanding of space systems' rigorous demands and challenges, which I keep porting to the space networking domain until today.

Then, through a competitive selection process, I won a concourse for a permanent researcher position at CONICET in Argentina, which I took in 2018. During this period, I undertook **multiple long-stay research visits** to prestigious institutions, including Politecnico di Torino in Italy, JPL (NASA) in the United States, Nanjing University in China, and Saarland University in Germany, where I remained as a guest professor within Prof. Holger Hermanns' chair. These collaborations, which remain active and are integral to the work presented in this HDR, significantly enriched my research and human skills and broadened my international academic network.

In 2021, I applied for and secured my current ISFP position at Inria, joining the Inria Agora team on the INSA Lyon campus, led by Prof. Hervé Rivano. This role allowed me to delve deeply into applying **computer science methods in satellite IoT**. I also secured enough funding to enrich and complete my research portfolio by integrating terrestrial and space-based communication systems, as summarized in this HDR.

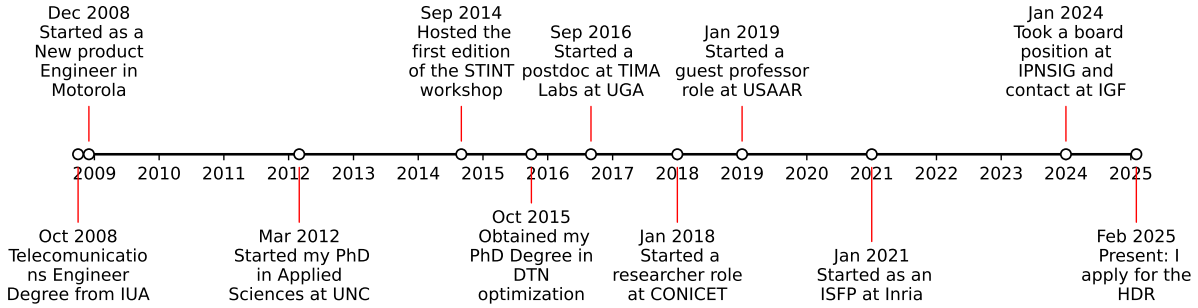

Figure 1.1: Trajectory timeline.

Memberships I am the founder and chair of the **Space-Terrestrial Internetworking Workshop** (STINT), a successful event since 2014. I co-chair this initiative alongside esteemed scientists such as Scott Burleigh, Edward Birrane, and other prominent figures in the field. Additionally, I serve as a board member of the **Interplanetary Networking Special Interest Group** (IPNSIG), a position I have held since early 2024. In this role, I engage in visionary discussions with Vint Cerf and other globally acclaimed Internet pioneers. I also contribute as an academic advisor to the Dynamic Coalition on the Interplanetary Internet at the **International Government Forum** (IGF).

Teaching Teaching has been a constant and fulfilling aspect of my career. While in Argentina, I lectured on **Computer Architectures and Networks** at FAMAF (UNC) and coordinated and taught various postgraduate courses at CONAE. I also developed and delivered my postgraduate course, **Space Informatics and Networking**, at both Politecnico di Torino and Saarland University, where I am a guest professor. Currently, I lecture on **Satellite Communications and Navigation** at INSA Lyon, collaborating with external experts from CNES.

Mentors Throughout my journey, I have been incredibly fortunate to have had exceptional mentors who provided invaluable guidance, knowledge, and inspiration. Their support has profoundly shaped my path and contributed much to who I am today. As I supervise my students now, I often find myself echoing actions and gestures I learned from them—a testament to their lasting impact on my approach to research and mentorship. For this, I will always be deeply grateful. My acknowledgment to them, today all close friends, is listed chronologically as follows.

- **Prof. Pablo A. Ferreyra** (UNC). Pablo taught me the importance of academic life's *human and emotional aspects*. He was my professor for C programming during my first year as a university student, and his guidance was instrumental in fostering

my passion for research to the point he became my PhD supervisor years later. I shared multiple joint projects involving CONAE in Argentina.

- **Prof. Jorge M. Finochietto** (UNC-CONICET). Jorge showed me that research can be *meaningful and impactful* and that high-quality questions can bridge the gap between academic rigor and practical applications. As my PhD co-supervisor, his mentorship was invaluable, and our collaboration continues to this day through ongoing projects and connections with students at UNC in Argentina.

- **Dr. Raoul Velazco** (TIMA Labs, CNRS, retired). Raoul extended a welcoming hand that brought me to Europe for the first time, serving as a *bridge to new opportunities*. His persistent humor and infectious positivity made my transition into European life, particularly in the charming city of Grenoble. I remain a close friend of Raoul and his family, to whom I will always be deeply grateful.

- **Scott Burleigh** (JPL, NASA, retired). Scott *deeply inspired my intellect beyond the horizon* with a profound, forward-looking motivation to explore the vast possibilities of space and interplanetary networks. His remarkable blend of literary insight and expertise in on-board software development left a lasting impression on me. Scott is the primary reason I am deeply engaged in space networking.

- **Prof. Holger Hermanns** (USAAR). Holger *infused the rigor and joy of computer science* into my career. Collaborating with him and his exceptional Depend team provided the pivotal bridge to computer science that dramatically propelled my career forward. From conducting in-orbit experiments to engaging in remarkable projects with ESA, I am deeply grateful to be a guest professor on Holger's team.

- **Prof. Hervé Rivano** (Inria-INSA Lyon). Hervé demonstrated that *ideal research conditions* can indeed exist. His unwavering support has significantly boosted my career over the past few years, providing me with the right tools, resources, and flexibility to sustain and expand the global collaborations that define my unique research profile at Inria today.

1.4 Publications

This section presents a high-level overview of my research contributions in the past decade, encompassing 120 published papers as of January 2025. All are referenced throughout this document and listed in the Bibliography section.

To maintain focus, clarity, and conciseness, I have chosen **not to include or cite related works by other authors** in this document. Instead, only my publications are

cited and discussed. Readers seeking additional context or background can access my works through their respective DOIs or HAL entries, which are conveniently provided. All cited papers have publicly accessible self-archival copies available online.

The remainder of this section organizes my contributions into distinct **Categories**. It provides a concise analysis of my research through the lenses of **Methodologies**, the **Roles** I took in these works, and the kind of **Venues** on which my work has been published.

1.4.1 Categories

The structure of my research is systematically organized into four distinct categories I discovered during the HDR process.

Figure 1.2 illustrates the distribution of my publications across these categories and their chronological evolution throughout my career. The plot illustrates an increase in productivity since I took my position at Inria, with a notable surge in 2024 as several ongoing projects began delivering outcomes, particularly in the links and access category, which aligns with my core integration project on Satellite IoT at Inria.

The categories, which serve as the foundation for the organization I use in the remainder of this HDR document, are outlined in detail below. They provide a structured framework for presenting my research contributions and aligning them with the thematic focus of the following chapters.

- **Applications, Architectures, and Environment**. This category addresses the fundamental questions: *Why use space?*, *How can space be utilized effectively?*, and *What are the characteristics of the space environment?* This approach provides readers with an intuitive framework to motivate the subsequent topics discussed in the remainder of the HDR manuscript. This category is addressed in Chapter 2 and detailed in Appendix A.

- **Analysis, Design, and Operations**. This category delves into the essential questions: *How can space networks be analyzed, designed, and operated effectively?* These contributions form a foundational framework for managing satellite constellations, designing robust contact plans, and optimizing mission operations. This category is summarized in Chapter 2 and deepened in Appendix B.

- **Routing and Flow Control**. This category revises contributions addressing the critical questions: *How to determine the next hop in multi-path satellite networks?* and *How to manage data flows effectively?* This chapter of the HDR bridges theoretical models with practical implementations, offering scalable solutions for multi-hop

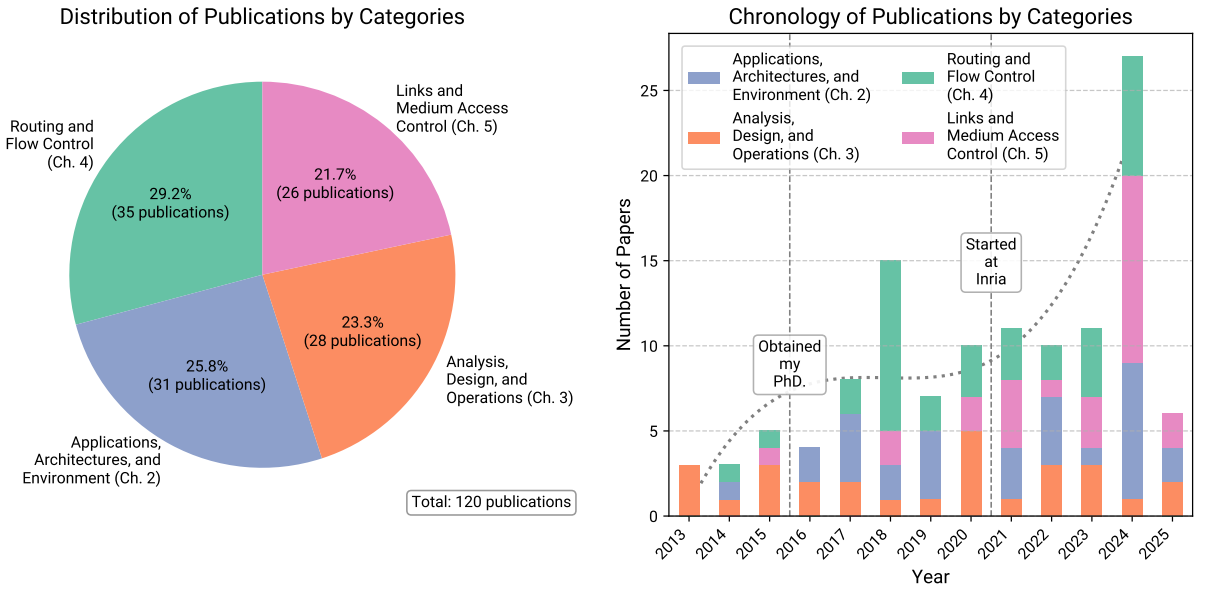

Figure 1.2: Publications by categories. **Left Panel**: The pie chart demonstrates a well-balanced effort across all four identified categories throughout my research trajectory. **Right Panel**: The chronological evolution of my publications by category. It shows an increase in the aggregate number of publications following my appointment at Inria. The plot also shows sustained productivity across other categories, reflecting a continuous effort supported by the collaborative environment fostered at Inria.

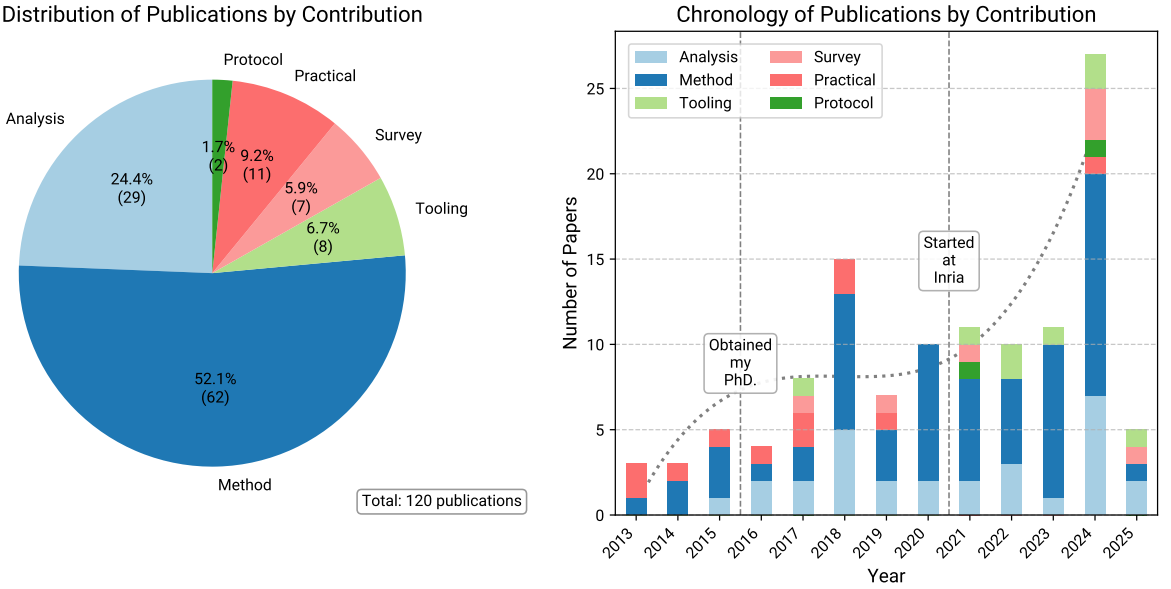

Figure 1.3: Publications by methodology. **Left Panel**: The pie chart highlights that most contributions focus on *methods* (e.g., optimization procedures, algorithms), followed by the *analysis* category. The practical category follows. It encompasses implementations and experimental validations. **Right Panel**: The chronological progression indicates a growing emphasis on developing methods, alongside a consistent focus on tooling and practical implementations, suggesting an evolution of my research towards applied contributions.

routing and transport in dynamic satellite networks. This category is overviewed in Chapter 2 and further elaborated in Appendix C.

- **Links and Medium Access Control**. This category dives into the fundamental question: *How can data transfer and multiple access be efficiently coordinated among devices across vast space distances?* This chapter lays the groundwork for understanding single-hop point-to-point communication challenges. This category is summarized in Chapter 2 and expanded in Appendix D.

1.4.2 Methodologies

Figure 1.3 illustrates the distribution and evolution of my research contributions based on their methodology. While many papers embrace a combination of methodologies, I tried to select the most predominant. I categorized the methodologies as follows:

- **Analysis**: Performance evaluations of specific space networks, typically via detailed simulations.
- **Method**: Proposals of new procedures or mechanisms, such as decision-making frameworks or optimization techniques.
- **Tooling**: Descriptions of software tools designed for community use, typically simulators or libraries, emphasizing usability and reproducibility.
- **Survey**: Comprehensive taxonomies of extensive research domains, often accompanied by tutorial-style explanations.
- **Practical**: Contributions involving protocol stack implementations or experimental validations, including in-orbit experiments.
- **Protocol**: Definition and specification of novel and ad-hoc communication protocols tailored to the unique needs of space networks.

These categories provide a comprehensive view of the methodologies driving my research contributions. The timeline in Figure 1.3 further highlights the evolution and balance of these methodologies over my career.

1.4.3 Roles

My role in publications has substantially evolved throughout my career, as depicted in Figure 1.4. Contributions are categorized into four distinct roles:

- **Lead**, where I serve as the primary author and undertake most of the effort.

- **Cooperation**, where the workload is collaboratively shared among multiple contributors.
- **Supervision**, where I act as the principal supervisor guiding a student who leads the paper[1].
- **Co-Supervision**, where I provide support as a secondary supervisor, working alongside another primary supervisor to mentor a student leading the publication.

This progression reflects my contributions' increasing scope and diversity, which have transitioned from individual-led efforts to a focus on mentorship and collaborative research.

My early research efforts during my PhD were primarily self-driven, resulting in publications where I took the lead. As my career advanced, **collaboration became a cornerstone of my work**, with an increasing number of co-authored publications focusing on teamwork and knowledge sharing. In recent years, most of my output has been driven by **supervised and co-supervised research** with Master's and PhD students.

1.4.4 Venues

Figure 1.5 expands on the analysis in the previous Figures by focusing on the type and quality of my publications throughout my career. Publications are categorized into **Books**, international **Conferences**, and *Journals*, with journals further classified into **Q1 Journals** (top-tier) and **Other Journals** (ranked in Q2 or lower quartiles).

The data reveal a consistent ≈3:2 ratio between conferences and journal publications, underscoring the balance between presenting at leading conferences and contributing to archival journals. Among journal publications, Q1 journals dominate with a strong ≈3:1 ratio over lower-ranked journals. Chronologically, the record highlights a significant increase in Q1 journal articles after joining Inria.

I am not particularly fond of metrics like journal impact factors or conference gradings. **Cross-disciplinary fields like space networks often struggle to find high-ranking venues that align with the niche scope**. Even when such venues exist, they may not always be the most attractive according to these academic metrics. That said, I acknowledge that the rigorous review processes in high-ranking venues significantly contribute

[1] In the supervision category, I include both directly supervised students (Santiago Henn, from CONICET) and co-supervised students for whom I assumed the primary supervisory role specifically for the given paper.

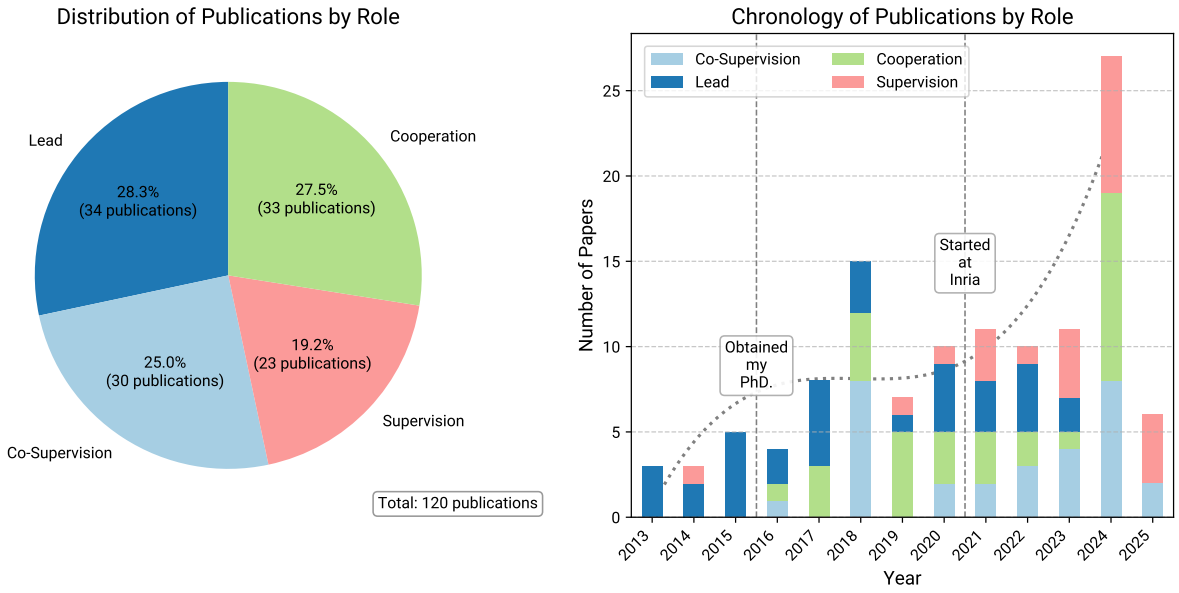

Figure 1.4: Publications by my role on them. **Left Panel**: The pie chart indicates that approximately half of my publications stem from supervised or co-supervised work with Master's and PhD students, highlighting my active mentorship role at this career stage. **Right Panel**: The chronological timeline reveals a clear progression in my publication roles. Early in my career, the majority of publications were self-led. Over time, this shifted to collaborations, and more recently, there has been a predominance of publications resulting from my supervision and co-supervision activities.

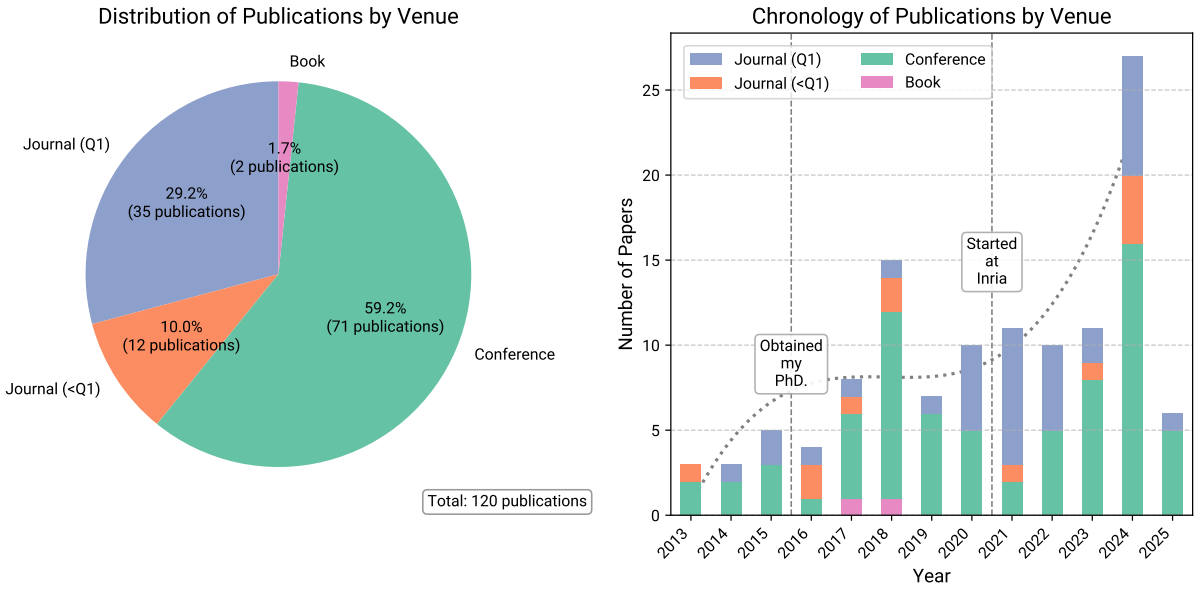

Figure 1.5: Publications by venue and quality. **Left Panel**: The pie chart shows the distribution of publications by type, showing a consistent 3:2 ratio between conference and journal papers. Among journal publications, there is a clear dominance of Q1 journals, with a 3:1 ratio compared to journals ranked below Q1. **Right Panel**: The chronological evolution highlights that most high-impact Q1 journal articles were published after joining Inria in 2021. A significant spike in conference publications is observed in 2024, reflecting the productive output of newly incorporated students under my supervision.

to the quality of the research output. In any case, I provide Figure 1.6, which lists my career's most frequent conference and journal venues.

Another important aspect regarding the quality index of publications is my **deliberate choice to target lower-impact journals for younger and less experienced students**. This approach is designed to help them navigate the complete publication cycle, including review rounds and proof stages, with manageable ideas. By doing so, I aim to build their confidence, foster learning, and, most importantly, minimize the risk of frustration during their early academic journey. This policy is reflected in the publication analysis shown in Figure 1.7, which presents and discusses the cross-relation between venue categories and role categories. The figure highlights how my strategic focus on impactful journal outputs aligns with my role as lead author. In contrast, other venues, such as conferences, predominantly feature contributions from co-supervisions and collaborations.

1.5 Funding

In recent years, I have secured funding to establish a highly productive research team, which explains the significant growth in publications highlighted in the previous section. Figure 1.8 presents a timeline of all active and concluded fundings discussed below.

1.5.1 Active Funding

I still benefit from funding obtained during my tenure at CONICET (e.g., the MISSION project), which supports missions to and from South America. Additionally, I manage funding acquired at Inria (e.g., the STEREO and DONUTS projects) and funding from ESA through my guest professor position at Saarland University. I also collaborated with two Italian universities, Politecnico di Torino and University of Genoa, to secure funding for hiring PhD candidates, with whom I co-supervise their research. A detailed list of secured and currently active funding is listed below.

- **STEREO** (Space-Terrestrial Integrated IoT): Funded by ANR PRCE (France), and with Inria as beneficiary, this project runs from March 2023 to March 2027 with a total funding of ≈€1M from which ≈€200k are available to our team. From these funds, we have hired Alexander Ylnner as a PhD candidate and Hannah Pasandi as a postdoc. As the Coordinator, I lead efforts to achieve an integrated space-terrestrial IoT architecture and protocols. The partnership includes various academic and industrial institutions from France: Inria (coordinator), IRIT, LAAS, LIG, and Kineis (industrial).

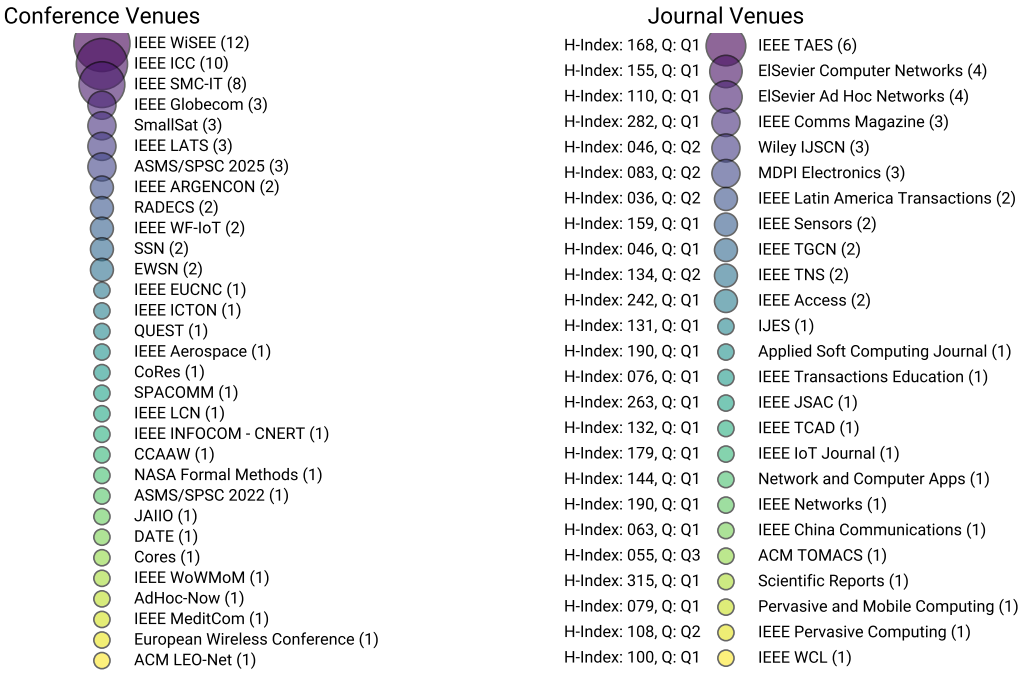

Figure 1.6: The distribution of my publications spans both conference and journal venues. On the conference side, venues such as IEEE WiSEE and IEEE SMC-IT stand out as pivotal intersections of space technologies and networking, where I often host my STINT workshop. On the journal side, esteemed publications like IEEE TAES, Elsevier Computer Networks, and Elsevier Ad Hoc Networks provide some of the most fitting Q1 platforms for advancing space networking research. H-index and Q rankings, obtained from scimagojr, are provided for reference.

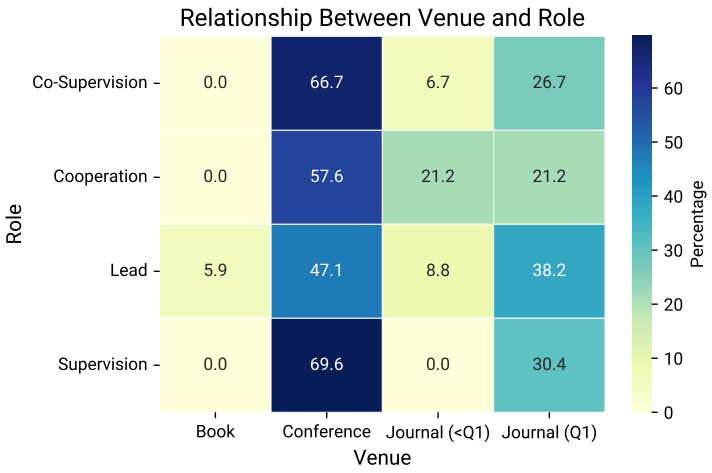

Figure 1.7: Figure Caption: Relationship between publication venues and roles. Most Q1 journal publications (38%) feature me as the first author, demonstrating a focus on high-impact outputs when leading the research. Conversely, I contributed the least to conference publications as the lead author, with these efforts primarily involving co-supervisions (66%) or collaborations (57%). Lower-tier journals (below Q1) make up only 8% of my lead-author contributions.

- **MISSION** (Models in Space Systems: Integration, Operation and Networking): This project, funded by EU RISE and with UNC and Saarland University as beneficiaries, runs from October 2021 to October 2026 (extended) and is funded at ≈€1.5M to support extensive secondments in Europe and Latin America. As a Responsible researcher, I contribute to developing scalable and dependable informatics models and methods for space missions. The partnership includes various academic and industrial institutions: UT, RWTH, USAAR, UNC, UNRC, D3TN (industrial), Ascentio (industrial), and INVAP (industrial).

- **DONUTS** (Design and Modeling of Multiscale Multitechnology Networks of the Future: From Cellular 5/6G to Non-Terrestrial Networks): This project, funded by the France 2030 program and with INSA Lyon as beneficiary, is scheduled to start in March 2025 and run until November 2029. It obtained funding of ≈€2.4M, from which ≈€200k will allow us to hire a PhD and a postdoc position on our team. As a project member, I will contribute to the space and terrestrial integration of satellite IoT. The partnership includes several universities and research institutes in France, including Université de Toulouse, Université de Strasbourg, ENS Lyon, INSA Lyon, Grenoble INP, LAAS, and CEA.

- **DORSAL-IOT** (Downlink Optimization for Robust Direct-to-Satellite LoRaWAN IoT): Funded by the Regional Program STIC-AmSud and with Inria as beneficiary, this project runs from January 2025 to January 2027, with a small funding to support missions between France and South America of ≈€33k. As a member of the coordinator group, I will lead efforts to design effective downlink techniques for LoRaWAN-based satellite IoT. The partnership of this project comprises: Inria (coordinator), Universidad de Chile, Universidad Diego Portales, UTEM, and UFSC from Brazil.

- **AI4LFNM** (Artificial Intelligence for Large Fleet Network Management): Funded by ESA and with Saarland University as beneficiary, this project runs from February 2023 to March 2025 with a total funding of ≈€500k, ≈€200k of which goes to the Saarland University which was partially used to support Gregory Stock PhD. As a lead Researcher, I focus on creating AI models to manage and operate space constellations. The project is led by GMV (industrial) with the Saarland University as a subcontractor and Thales Alenia Space, Eutelsat, and Planet Labs as partners.

- **ML4OISL** (Machine Learning for Optical Inter-Satellite Link): Funded by ESA with also Saarland University as one of the beneficiaries, this project runs from March 2024 to October 2025 with a total funding of ≈€500k, ≈€75k of which goes to the Saarland University which was used to hire Belen Llaneza, a support engineer I

supervise. The partnership includes GMV (industrial) as the prime contractor and Saarland University and Eutelsat as subcontractors.

- **HANDING-OVER** (Handover, Data Routing, and Radio Resource Management for Very Low-Earth Orbit (VLEO) Broadband Constellations): Funded by ESA I supported Politecnico di Torino as a beneficiary. The project is funded with ≈€500k, part of which was used to support Gabriel Maiolini Capez, a PhD student I co-supervise. The project's prime contractor is Thales Alenia Space, Italy.

- **PhD Scolarships**: I currently fund three of my PhD students with dedicated scholarships as follows.

  - **ED Lyon PhD Scholarship**: I secured funds for a full PhD The BMI scholarship from Lyon Ecole Doctorale (InfoMath) was for Diego Maldonado, who started in December 2022 and is scheduled to finish in December 2025.
  - **CONICET PhD Scolarship**: Before taking my Inria ISFP position, I secured funds for a full PhD position awarded to Santiago Henn. Starting in December 2022, he is scheduled to defend also in December 2025[2].
  - **UNIGE PhD Scolarship**: I supported Fabio Patrone from University of Genoa to find funds and co-supervise Camilo Rojas. He started in January 2023 and is scheduled to defend in January 2027.

1.5.2 Concluded Funding

I have also benefited from additional funding sources that have since concluded. The list below provided the economic resources for my two-year research period at CONICET between 2018 and 2020 and partially overlapped with my first period as an ISFP researcher at Inria.

- **i2Cat Contract** (Direct collaboration between i2Cat and Inria): This contract, funded by i2CAT in Barcelona, Spain, provided approximately €50k and spanned the entire year of 2023, with funds available through the end of 2024, used mostly for conferences, mobility, and equipment. The collaboration focused on advancing research in optical inter-satellite links and network management. A key outcome of this partnership is my ongoing co-supervision of Elena Fernandez-Niño and Arnau Singla for their PhD projects. Despite the contract's expiration, the collaboration remains active, and the supervisory role continues to foster impactful research in the field.

[2]Because of Santiago Henn's supervision and the continuity of the MISSION project, I remain part of CONICET, although detached from my position (i.e., temporary leave).

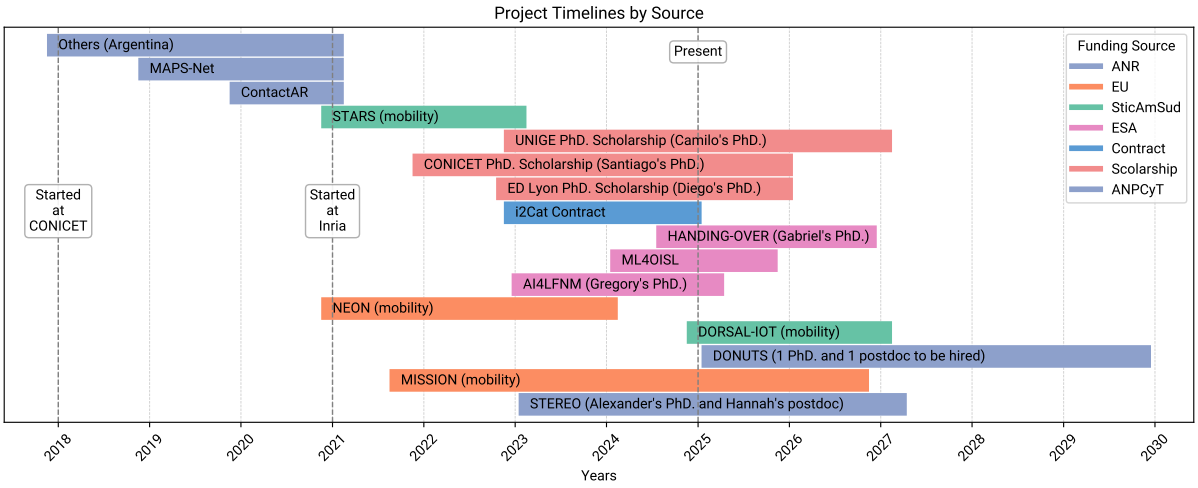

Figure 1.8: Funding timelines by source. This figure illustrates the timeline of my funded projects categorized by their funding sources. The timeline begins with my tenure at CONICET in Argentina, showcasing early projects funded by ANPCyT. It then transitions to my current role at Inria, marked by an expanded portfolio of projects such as NEON (EU funding), STARS (SticAmSud mobility), and ESA-backed initiatives like HANDING-OVER. The chart highlights the sustained productivity and diversity of funding sources, including ANR, EU, SticAmSud, ESA, scholarships, and ANPCyT, which have supported PhD students, postdocs, and mobility programs.

- **NEON**: Funded by the Erasmus+ Program (EU), this project runs from 2021 to 2024 with total funding to support international courses and seminars of ≈€1M. As the Responsible Professor, I contribute to IoT capacity building in ground and space domains, making it relevant to AsyncNet. The partnership of this project included: UCM3, UNC, UNMdP, UNS, and others.

- **STARS** (Satellite Networks Architectures Protocols and Informatics): A SticAmSud Project, that ran from 2021 to 2023, with UNC, Universidad de Chile, and LAAS as partners. This project supported missions between South America and France before and during my role at Inria with a funding of ≈€30k.

- **ContactAR** (Contact Traceability through the Digital Context of Mobile Devices): A project funded by ANPCyT to support research related to the COVID-19 pandemic from 2020 to 2021 with the equivalent of ≈€150k.

- **MAPS-Net** (Models Protocols and Algorithms for Segmented Satellite Networks): This project funded by ANPCyT is analogous to ANR's JCJC but in Argentina. It funded my early research from 2019 to 2021 with the equivalent of ≈€60k.

- **Others**: Projects such as "Low Power Wide Area Communication Models for the Internet of Things", "Proactive controllers for software-defined networks", and "Trans-

port Networks Reliability: Models, Algorithms and Protocols", all funded by ANPCyT, provided additional small but continuous flow of money to conduct research between 2018 and 2021.

1.6 Group

While my institutional affiliation is with Inria Agora, where I actively collaborate with many permanent members, I lead a virtual research network that I informally refer to as "Space Informatics and Networking" (SPAIN), an acronym I must credit to Prof. Holger Hermanns.

My research network is unique in its composition. It brings together researchers and students from diverse institutions and fosters a highly interdisciplinary and collaborative environment across engineering, computer science, and international boundaries. Such diversity, possibly a consequence of my "nomad" career, is a strength and a necessity in the space context. As I discuss in Chapter 3, **any real space network must embody this spirit of cooperation, transcending national frontiers**—because, in space, there are no borders, at least not yet.

The cross-border setup of my research group proves immensely beneficial in gathering diverse space networking perspectives from space agencies worldwide. By engaging with researchers and institutions across continents, the team gains valuable insights into different space programs' priorities, challenges, and strategies. This global outlook enriches the research and fosters collaborations that align with space's inherently borderless nature.

Figure 1.9 provides a visual representation of the structure of SPAIN. It illustrates the breadth of my professional research network, categorizing key relationships into mentors, directly supervised students, co-supervised students, and collaborators. Each individual in this figure is referenced at least once in the subsequent chapters of this HDR, reflecting their integral role in my research journey.

Mentors I have already highlighted the pivotal roles played by Prof. Hervé Rivano, Prof. Holger Hermanns, Scott Burleigh, Dr. Raoul Velazco, Prof. Jorge M. Finochietto, and Prof. Pablo A. Ferreyra in shaping and guiding my career. Their mentorship has been instrumental in both my academic and professional development.

Directly Supervised Students I have one directly supervised PhD student, Santiago Henn, who remains under my sole supervision, reflecting the authority I held during my

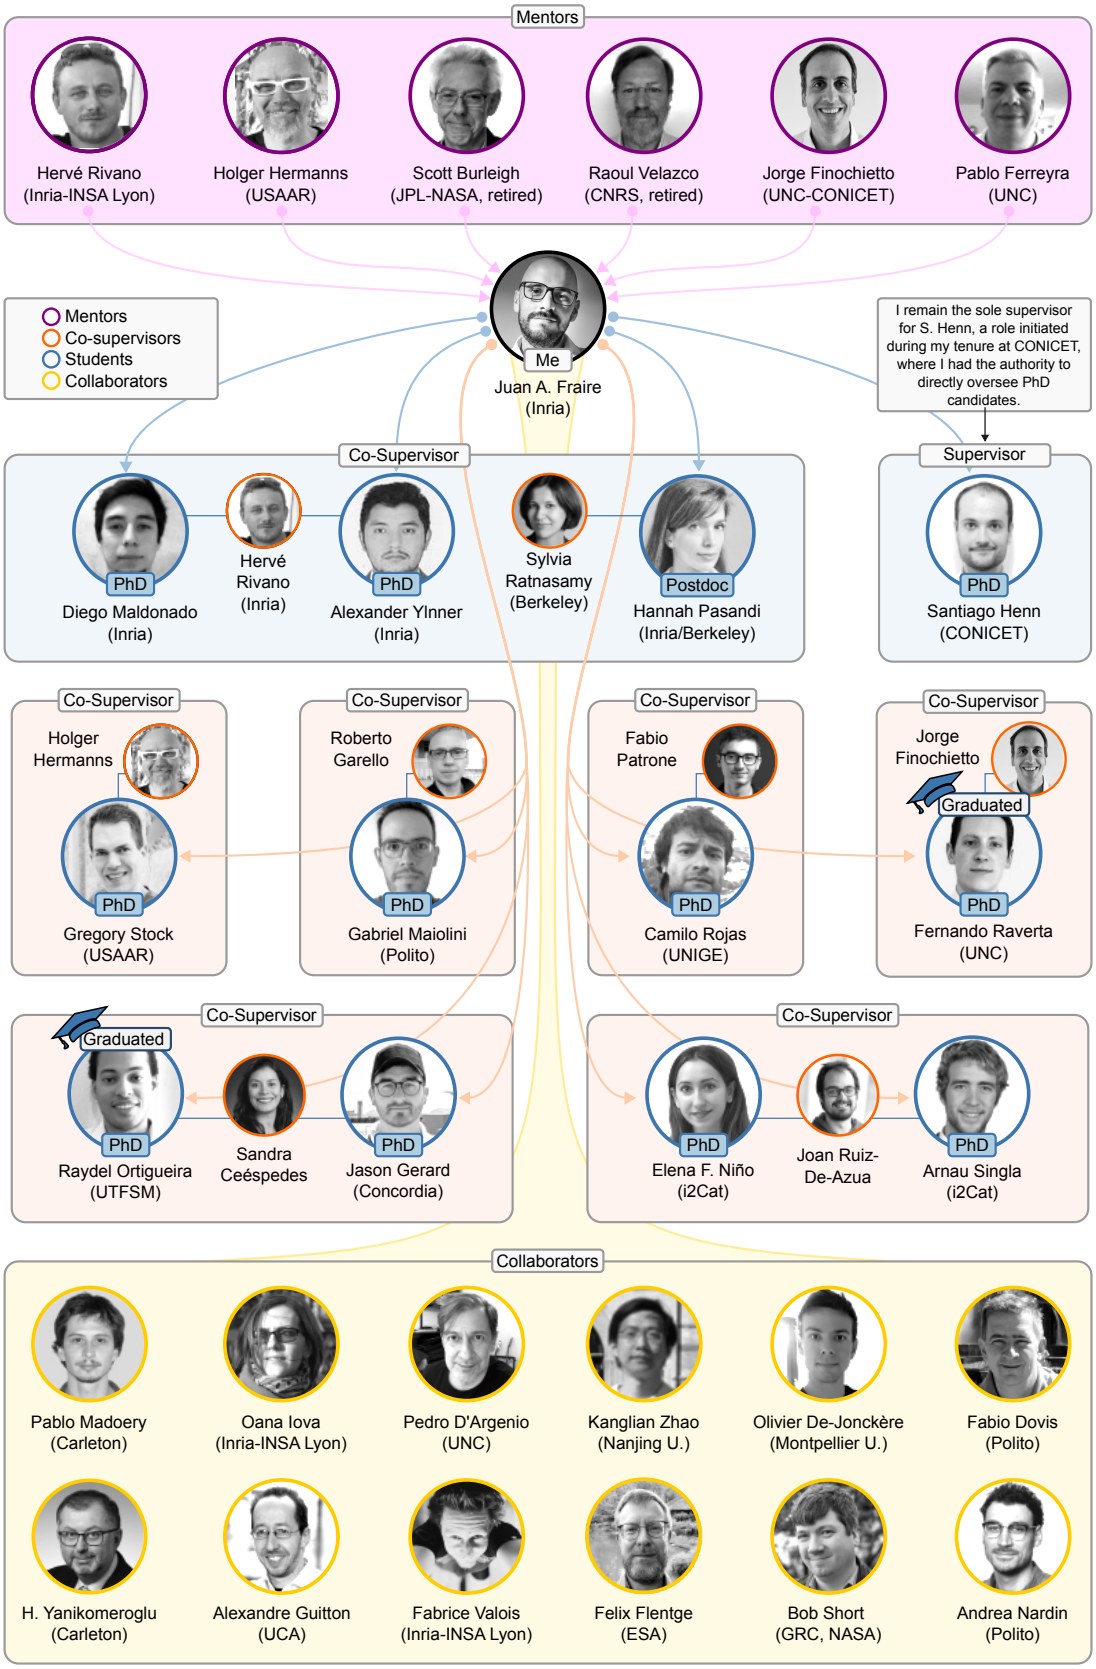

Figure 1.9: The “Space Informatics and Networking” (SPAIN) research network.

tenure at CONICET, where I was allowed to directly supervise PhD candidates (while the HDR will grant me this privilege in France).

Co-Supervised Students The figure highlights the students whose research I co-supervise in collaboration with esteemed colleagues across various institutions. Among them, Fernando Raverta (from UNC) and Raydel Ortigueira (from UC Berkeley) have completed their PhD studies. Both were co-supervised, with myself serving as a formally recognized advisor alongside their primary supervisors.

Diego Maldonado and Alexander Ylnner, both PhD candidates, are integral members of my core team at Inria Agora (co-supervised with Prof. Hervé Rivano at Inria, where I have taken the lead role on a day-to-day basis). I also co-supervise Hannah Pasandi, a part-time postdoc at our team working with Sylvia Ratnasamy from UC Berkeley.

The remaining students, represented in orange in Figure 1.9, are under the primary guidance of other colleagues; my role with them is that of secondary co-supervisor. Among these, I co-supervise Gregory Stock (from Saarland University), Gabriel Maiolini Capez (from Politecnico di Torino), Camilo Rojas (from University of Genoa), Jason Gerard (from Concordia University), and Elena Fernandez-Niño and Arnau Singla (both from i2CAT) in collaboration with co-supervisors from several institutions, fostering cross-disciplinary and multi-institutional research initiatives. The high publication throughput detailed in this HDR stems from the collaborative efforts within this extensive and active research network.

Collaborators The collaborators listed include prominent engineers, researchers, and decision-makers with whom I have collaborated impactfully. This shows the interdisciplinary nature of my research in space informatics and networking. The collaborators listed in the figure are the most frequent co-authors of my papers.

Master's Students Furthermore, numerous internships and Master's students have contributed to this work, such as Marcos Alfonzo, Blas Vega, and Matias Olmedo (the three from CONAE); Elias Gasparini and Dardo Ariel Viñas (from UNC); Mehdi Ait Mesbah, Pablo Illabaca, Sebastian Montoya, Benoit Coeugnet, and Alice Le-Bihan (the five of them from Inria); Carsten Gerstacker, Kai Vogelgesang, and Robin Ohs (the three from Saarland University), and Camilla Ottaviani (from Politecnico di Torino), all of which managed to deliver works that where later accepted to publication in conferences and journals. Their efforts and contributions are acknowledged and referenced throughout this HDR. Master's students whose work was not published are not listed in this document.

Industrial Partners Last but not least, I would like to highlight the many companies I collaborate with across diverse domains, including GomSpace for satellite operations, D3TN GmbH for delay-tolerant networking protocols, Sateliot, Kinéis, and Pontos for direct-to-satellite IoT, and Skyloom for free-space optical networks. I refer to these companies in multiple parts of the HDR. I interact with all of these industrial partners in the context of the MISSION or STEREO projects.

> **Industrial partnerships** have kept me talking the "space industry" language, bridging research with real-world solutions while advancing the state of the art in their respective application fields.

I began my journey as an engineer, gaining industry expertise at Motorola, where I learned the language of the field. My PhD was rooted in applied computer science, gradually leading me deeper into theoretical computer science. Currently, I work and collaborate with distinguished computer scientists. **This unique trajectory places me in a distinctive interdisciplinary role—bridging both satellites and the gap between theoretical and applied science**.

> **Informatics as a Strategic Gap**. Through my collaborations with industrial partners in the space networking sector, I have identified a critical and often overlooked gap: **the lack of advanced informatics models and methods tailored to address the space industry's unique challenges**. The space industry, deeply rooted in technically driven engineering processes, operates under strict protocols and rigid workflows designed to meet the high-stakes demands of its politically sensitive and harsh operating environment. **This leaves little room for exploring abstract or forward-looking methodologies** that fall outside their immediate scope.

My research aims to bridge this gap by providing these much-needed models and methods through direct collaborations or by training a new generation of highly skilled PhD students who can integrate into these companies and drive innovation from within. I elaborate on my forward-looking vision in greater detail in Chapter 3.

1.7 Reflection

This is the ideal time for me to apply for the HDR. The following core aspects highlight this decision's motivations and key reflections

Over the years, I have successfully guided two PhD students to graduation and co-supervised several others, gaining invaluable insights from collaborating with mentors, institutions, and cultures worldwide. These experiences have equipped me with the confidence and capability to independently lead research and articulate my authentic vision for the field. Even during periods of sole supervision, I consistently delivered meaningful and impactful contributions. Furthermore, I have secured competitive funding as a researcher and coordinator of complex, multi-institutional projects which bridges academic and industrial partners. These achievements demonstrate my readiness to formalize my leadership in space networking and inspire the next generation of researchers.

Chapter 2

Models and Methods

This chapter consolidates my research contributions at the intersection of computer science and the space engineering domain. I employ **advanced models and methods to address the unique challenges of networked space systems**.

These contributions span theoretical frameworks, practical implementations, and experimental validations. They leverage innovative techniques to optimize applications, architectures, routing, and medium access control in space networks.

The chapter aligns with the research categories introduced in Section 1.4.1: *i*) **Applications, Architectures, and Environment**, *ii*) **Analysis, Design, and Operations**, *iii*) **Routing and Flow Control**, and *iv*) **Links and Medium Access Control**.

The appendices A to D, thoughtfully referenced throughout this chapter and aligned with its structure, provide detailed figures and optional supplementary materials that comprehensively support the works discussed. They provide additional context and validation for the methods and results discussed herein. Figure 2.1 depicts the context of the remainder of this chapter.

2.1 Applications, Architectures, and Environment

This section highlights my research contributions addressing fundamental questions in space exploration and networking: *Why use space?*, *How can space be utilized effectively?*, and *What are the characteristics of the space environment?*

My work focuses on the transformative potential of space-based services. It explores applications and architectures that extend terrestrial connectivity, enable robust interplanetary networks, and adapt communication paradigms to space constraints. Through simulation, experimentation, and innovation, I have advanced novel **Applications** and **Architectures**, redefining terrestrial and extraterrestrial communication systems. This section is further expanded in Appendix A.

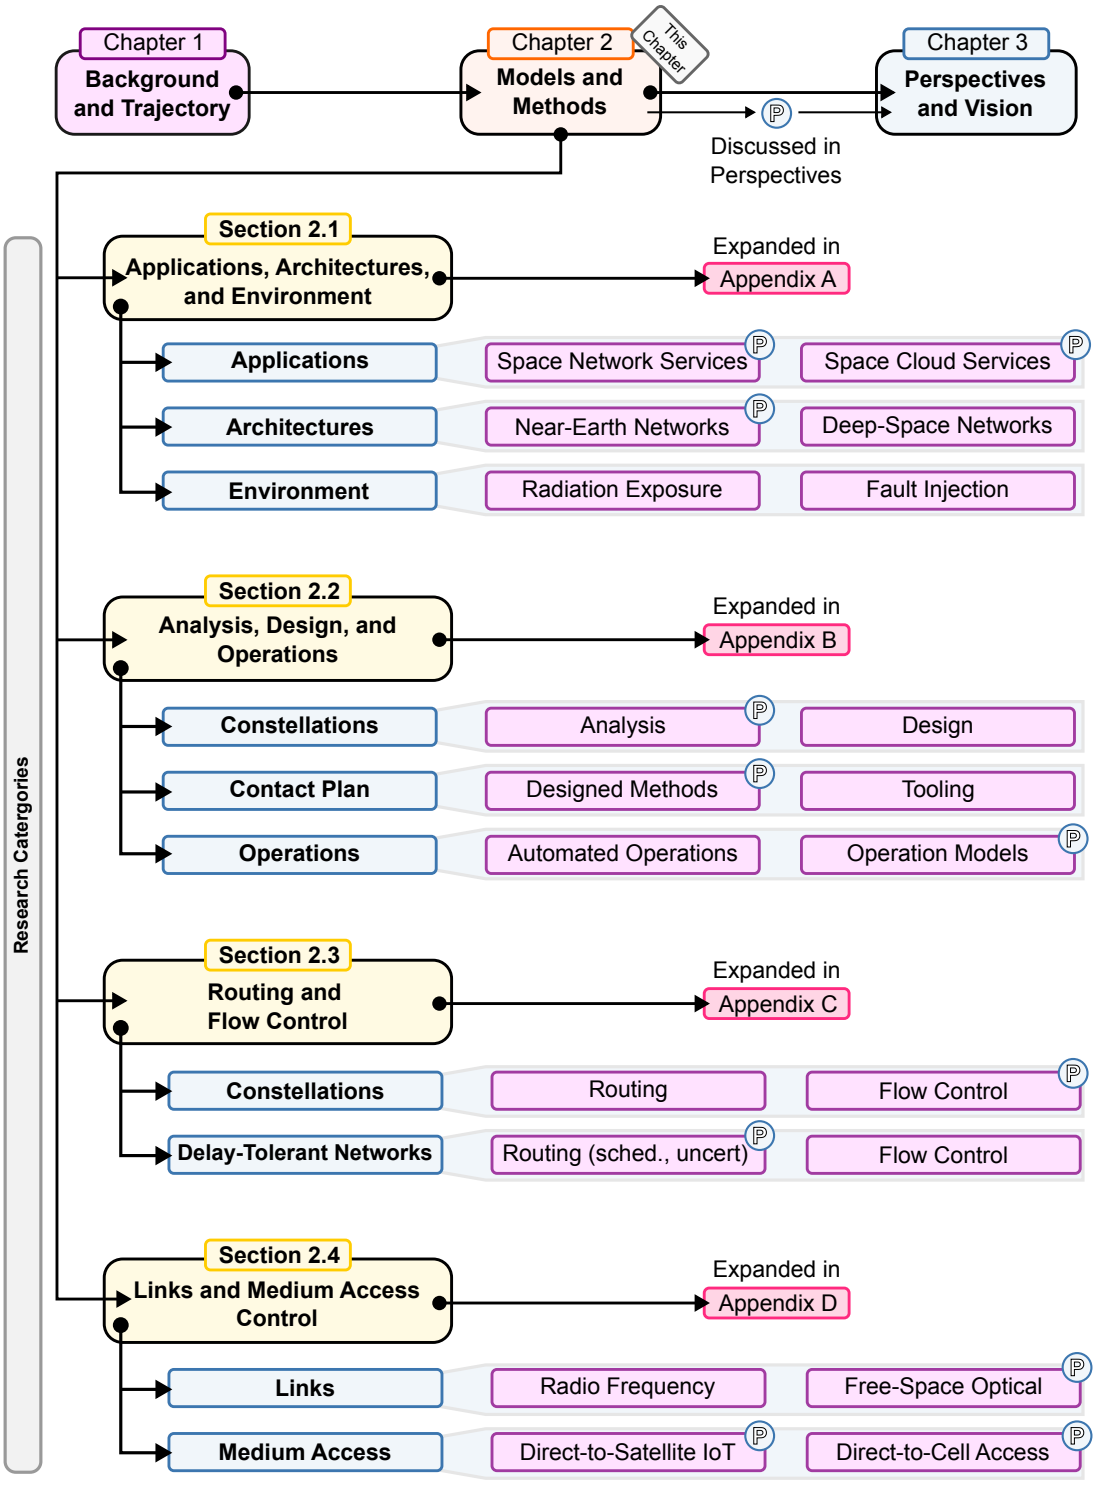

Figure 2.1: Outline of the Models and Methods Chapter.

2.1.1 Applications

This section examines the fundamental question: *Why use space?* Through rigorous research and experimental validation, it presents models and methods for two pivotal domains of space applications at the core of my research scope: **Space Network Services** and **Space Cloud Services**. The former facilitates data transport across satellite infrastructures, extending connectivity to underserved regions. At the same time, the latter enables local data processing within space assets, reducing latency for delay-sensitive tasks from the end-user perspective.

2.1.1.1 Space Network Services

Space Network Services enhances connectivity by leveraging satellites to transport **observation** and **user** data to and from remote areas. My core contributions essentially contribute performance evaluation for concrete mission architectures and include:

- **Observation Data Transport**: (More insights in A.1.1.1.)
  - **Optical Observation Data Handling via DTN**: Optical Earth observation data handling using Delay-Tolerant Networking (DTN) introduced efficient segmented satellite constellations with opportunistic Inter-Satellite Links (ISL) was studied in early work [MFF16].
  - **S-AIS for the PANSAT Mission**: A study on Satellite Automatic Identification System (S-AIS) for the PANSAT mission, in cooperation with CONAE, evaluated antennas, attitudes, and data processing, directly informing mission design [Alv+21]. (More in Figure A.1.)
  - **SARE Analysis in MISSION**: Analysis of SARE missions for Argentina's maritime zones in the context of the MISSION project demonstrated the effectiveness of DTN-based transmission efficiency for uniform coverage and optimal revisit times [Fra+24]. (More in Figures A.2 and A.3.)
- **User Data Transport**: (More insights in A.1.1.2.)
  - **Ring Road Networks**: In cooperation with D3TN GmbH, I flight validated the Ring Road Network (RRN) concept on ESA's OPS-SAT mission, showcasing asynchronous connectivity using the μD3TN protocol, with applications extending to delay-tolerant web browsing, remote machine-learning services, and interoperable DTN experiments [Fel+22]. This marked the first in-flight test of Bundle Protocol (BP) Version 7 (BPv7) and cross-implementation interoperability. (More in Figure A.4, Section A.1.1.2.)

- **Underwater to Space Communications**: The ColdSun experiment demonstrated underwater-to-satellite data communication using acoustic underwater links and commercial satellite networks, achieving stable transmission under dynamic conditions, paving the way for cost-efficient maritime IoT systems. (More in Figure A.5, Section A.1.1.2.) I plan to deepen in this domain as discussed in Chapter 3.
- **Cross-Agency Communications**: Experiments in the context of the Lunar IceCube mission successfully validating BP interoperability across multiple DTN implementations in a simulated lunar mission scenario. The primary achievement of this experiment was administrative, successfully coordinating collaboration among NASA, ESA, and multiple private companies to prove a shared space-terrestrial network is possible. (More in Figure A.6, Section A.1.1.2.)
- **Mega-Constellation Services in Space**: The Mega-Constellation Services in Space (MCSS) paradigm introduced a scalable approach for integrating LEO satellites into 6G NTNs mega-constellations at higher altitudes, unlocking an appealing new ground-station-less paradigm [Cap+24]. (More in Figure A.7.)

Key DTN Experiments. In collaboration with D3TN GmbH, I spearheaded groundbreaking in-orbit, ground-to-space, and underwater-to-space communication experiments. These include: **a)** demonstrating the asynchronous connectivity of the μD3TN protocol on ESA's OPS-SAT mission; **b)** the ColdSun experiment, showcasing cost-effective underwater-to-satellite data transmission; and **c)** the Lunar IceCube Cross-Agency Experiment, which established a benchmark for cross-agency space communication. This underscores my initial assertion that the **methods and models in my research portfolio are deeply rooted in practical applications**.

2.1.1.2 Space Cloud Services

Space Cloud Services addresses service latency challenges by enabling in-orbit distributed processing. Camilo Rojas (University of Genoa) drove contributions in this domain. They focus on:

- **Server Orchestration**: We developed machine learning-based mechanisms, such as federated learning and reinforcement learning, for energy-efficient activation of edge servers onboard LEO satellites in a constellation, achieving a balance between performance and cost [Roj+24b]. (More in Figure A.8.)

- **Server Placement**: We analyzed latency trade-offs for different server placements in the constellation, revealing that space cloud configurations outperform the use of terrestrial clouds for globally distributed tasks, especially at distances exceeding 5000 km from the gateway [Roj+25]. (More in Figures A.9 and A.10.)
- **Tooling: MeteorNet**: The MeteorNet[1] emulation tool provides a robust containerized platform for validating edge computing strategies and generating synthetic datasets, bridging theoretical and practical developments in NTNs [Roj+23; Roj+24a].

Space Cloud. Space cloud services redefine NTNs by enabling on-satellite computation, eliminating dependency on terrestrial internet infrastructure. This innovation drastically reduces latency for delay-sensitive applications, positioning itself as a **game changer for 6G and beyond**, where NTNs are the driving use case. Addressing Space Cloud challenges forms a key perspective I further discuss in Chapter 3.

2.1.2 Architectures

This section addresses the critical question: *How to use space?* It re-evaluates conventional Internet architectures, providing models and methods to adapt them to the delays and disruptions inherent to space environments. Contributions span two domains: **Near-Earth Networks**, which optimize existing protocols for satellite systems below geostationary orbit, and **Deep-Space Networks**, which develop architectures capable of spanning the Solar System and beyond.

2.1.2.1 Near-Earth Networks

Near-Earth Networks adapt terrestrial technologies to deal with the dynamic connectivity and ranges of Low Earth Orbit (LEO). Key contributions include:

- **High-Altitude Platform Stations (HAPS)**: In cooperation with Carleton University, we proposed routing solutions to integrate HAPS into all-optical satellite constellations. These methods optimize data delivery by mitigating Free-Space Optical (FSO) link disruptions due to adverse weather, achieving higher delivery ratios and lower delays [Mad+24]. This novel HAPS-based architecture eliminates the need for Radio Frequency (RF) backups, offering transformative reliability enhancements for delay-tolerant LEO networks. (More in Figure A.11.)
- **Direct-to-Satellite Internet of Things (DtS-IoT)**: (More insights in A.2.1.2.)

[1] MeteorNet public repository: https://gitlab.com/camilo.rojas/meteornet

  - **DtS-IoT Survey**: I defined and formalized the DtS-IoT concept early in [FCA19], highlighting its unique challenges, such as long distances, Doppler effects, and energy efficiency. While initially met with skepticism from colleagues, DtS-IoT has now become a reality, not only in academia but also in the industry. (More in Figure A.12.)
  - **DtS-IoT Architecture**: Introduced in [FIV22], the Space-Terrestrial Integrated Internet of Things (STEREO) architecture integrates terrestrial and satellite IoT networks, focusing on direct rather than indirect satellite access. This work laid the foundation for my ANR-funded STEREO project, addressing key challenges in network function virtualization and mobility management. (More in Figure A.13.)
  - **Tooling: FLoRaSat**: I developed FLoRaSat[2], a simulation tool for DtS-IoT systems. This tool integrates orbital dynamics, ISLs, and LoRaWAN protocols, enabling end-to-end performance analysis under realistic scenarios [Fra+22b; Fra+22a]. The advanced FLoRaSat 2 supports large-scale LEO constellations with dynamic topology control, enabling comparative evaluations of routing algorithms like DiscoRoute [Cho+25]. (More in Figures A.14 and A.15.)

- **Educational Impact**: Based on the expertise on tools like FLoRaSat, I developed educational methodologies integrating simulation frameworks like OMNeT++ into networking curricula, enhancing student engagement and technical proficiency [FD21].

DtS-IoT. As the originator of the term DtS-IoT, I pioneered the concept of leveraging low-data rate, low-volume, and latency-tolerant applications to **not only relax the traditionally stringent requirements of space platforms but also optimize energy efficiency for battery-powered ground devices**. This paradigm formed the foundation of the integration plan at Inria, enabling cost-effective, scalable IoT solutions tailored for remote and underserved regions.

2.1.2.2 Deep-Space Networks

Deep-Space Networks address disruptions caused by planetary occlusions and extreme delays over vast distances. These effects essentially break any synchronous operation concept based on the current Internet. In this context, my contributions to deep-space networks focus on:

- **Interplanetary Networking**: (More insights in A.2.2.)

[2]FLoRaSat public repository: https://gitlab.inria.fr/jfraire/florasat

  - **ESA's Pre-Phase A Solar System Internet**: In cooperation with D3TN GmbH, I actively participate in ESA 's Solar System Internet (SSI) project, promoting scalable, autonomous, and standardized network designs for interplanetary contexts [FFB17a]. (More in Figure A.16.)
  - **Mars DtS-IoT Architecture**: I designed forward-looking architectures for Mars DtS-IoT services, integrating DTN and LoRa technologies for surface sensors, orbiters, and Earth communication, considering Martian environmental constraints [Led+24]. (More in Figure A.17.) As discussed in Chapter 3, I am convinced that *interplanetary IoT offers a transformative approach to space exploration*, complementing traditional rover-based methods with new broader coverage capabilities.

- **Tooling: IPN-V**: Supported by Alice Le-Bihan, I developed IPN-V[3], a visualization and planning tool for interplanetary DTN scenarios. The tool bridges astrodynamics and network engineering, offering interactive 3D visualizations of network dynamics [Le +24]. The graphical output of the tool is currently being used by ESA and D3TN GmbH in SSI related projects, and it was integrated into VR headsets and NASA's DTN toolchains [Bul+24]. (More in Figure A.18.)

- **Protocol Stacks**:

  - I worked in extending NASA's ION software with a SpaceWire convergence-layer adapter, enabling DTN protocol compatibility for intra-spacecraft communication [Alf+14]. This code was ported and used in CONAE's SARE mission and integrated with the Proximity-1 model discussed in Section 2.4. (More in Figure A.19.)
  - I collaborated on developing, testing, and deploying μD3TN, an open-source modular DTN protocol stack with advanced features like bundle encapsulation and scalable routing mechanisms. Validated through ESA's OPS-SAT mission, μD3TN addresses diverse terrestrial and space network needs [Wal+24]. (More in Figure A.20.)

- **Interstellar Networking**: I investigated interstellar communication feasibility using DTN principles for networks spanning light-years, addressing challenges like relay spacing, buffer size, and delay optimization for systems like TRAPPIST-1 [Fra+19]. (More in Figure A.21.)

[3] IPN-V public repository: https://gitlab.inria.fr/jfraire/ipn-v

Interplanetary Networks. My PhD research on CONAE 's SARE mission, with its sparse topology and episodic ISL communications, provided a foundation for addressing **analogous challenges in deep-space networks**. These networks share dynamics like planetary occlusion, extreme delays, and resource constraints, but they require distinct architectures compared to traditional constellations. My contributions naturally extended to interplanetary and solar system-scale networks, advancing the design and operation of deep-space communication systems.

2.1.3 Environment

This section examines the impact of the harsh space environment on electronic systems, with a particular focus on radiation effects, including Single-Event Effects (SEE) and Total Ionizing Dose (TID), which directly threaten the integrity of networking and communication systems. These effects pose significant challenges for satellite networks, disrupting modems, transponders, and other critical components.

This research advances the understanding and mitigation of radiation-induced failures, contributing to more reliable and robust space communication architectures. Contributions in this domain include **Radiation Exposure Campaigns**, **Fault Injection Campaigns**, and **Applications in Networks on Chip**. Most of them correspond to my stay at TIMA Laboratories for my postdoc.

2.1.3.1 Radiation Exposure Campaigns

Radiation exposure campaigns evaluate electronic resilience under controlled conditions using particle accelerators or heavy-ion beams. Contributions include:

- **Wireless Access for Radiation Beam Experiments**: I supervised the development of a novel wireless system integrating Bluetooth-enabled microcontrollers to enable real-time data acquisition during radiation tests. This approach, validated at the TANDAR facility, Argentina, improves safety, scalability, and reliability in testing components like SRAM for radiation-induced effects [Sán+19].
- **Memory Characterization**: I conducted static irradiation tests on SRAM devices at varying voltages, revealing increased sensitivity to SEUs at ultra-low biases and the efficacy of error-correcting codes (ECC) in mitigating multi-bit upsets and single-event functional interrupts [Cle+17; Cle+18]. The radiation campaign was executed in collaboration with Juan Clemente (UCM, Spain) using the GENEPI accelerator facilities in Grenoble, France.

2.1.3.2 Fault Injection Campaigns

Fault injection simulates radiation-induced errors in electronic systems to assess their fault tolerance. Contributions include:

- **NETFI-2 Framework**: Introduced a novel fault injection methodology for FPGA-based systems, leveraging an embedded controller for real-time fault injection. This method achieved faster execution, reduced resource overhead, and improved scalability while maintaining high accuracy in testing SEU and SET vulnerabilities [Sol+17]. (More in Figure A.22.)
- **Characterization of Bayesian Machines**: Assessed the robustness of stochastic Bayesian Machines (SBM), demonstrating fault tolerance under extreme fault rates, validating their suitability for probabilistic applications in radiation-prone environments [Coe+16; Coe+17]. This work was held in collaboration with Probayes SAS, a company from Grenoble, France.

2.1.3.3 Applications in Networks on Chip

Research extended to Networks on Chip (NoC) applications to improve their resilience to radiation effects:

- **Delay-Tolerant NoCs**: Adapted DTN principles to chip-level communication, integrating store-carry-forward mechanisms. Conceptual simulation results showed significant improvements in reliability, availability, and fault tolerance compared to traditional TCP-like NoC [Fer+19; Fer+21].
- **Route Computation Units for 3D NoCs**: Developed a fault detection and rerouting strategy to enhance the resilience of 3D NoC route computation units, reducing error rates and eliminating reliance on retransmissions with minimal area and power overhead. The module was implemented in VHDL and evaluated using realistic executions in FGPAs [Coe+18].

Evolution in the Field of Applications, Architectures, and Environment: I introduced **scalable and cost-efficient data transport and processing models**, exemplified by validating DTN protocols on ESA's OPS-SAT mission and the groundbreaking ColdSun underwater-to-satellite communication experiment. In addition, I proposed novel **architectural definitions**, contributing to the STEREO framework for DtS-IoT and novel server orchestration strategies in space cloud services. My work has bridged the gap between terrestrial and extraterrestrial

communication systems by addressing the challenges of latency, resource constraints, and **environmental resilience**. It has driven connectivity, computational efficiency, and network reliability advancements for near-Earth and deep-space applications.

2.2 Analysis, Design, and Operations

This section addresses the core question: *How can space networks be analyzed, designed, and operated effectively?* It covers emerging LEO satellite constellations, the topologies defined by contact plans, and the critical resource management (e.g., energy and memory) necessary for efficient operation.

Contributions are organized into **Constellation Analysis and Design**, **Contact Plan Design**, and **Mission Operations**, providing a comprehensive framework for ensuring the viability and optimization of space networks. This section is further expanded in Appendix B.

2.2.1 Constellation Analysis and Design

Satellite constellations are thriving (see Figure B.1). This section explores satellite constellations' dynamics and optimization, addressing the challenges of global and regional connectivity. It spans **Constellation Analysis**, which focuses on operational dynamics and performance assessment, and **Constellation Design**, which optimizes *sparse* configurations for delay-tolerant applications. These contributions emphasize enhancing coverage, scalability, and cost-efficiency while addressing challenges specific to dense and sparse constellations.

2.2.1.1 Constellation Analysis

Satellite constellations are intricate systems operating over vast spatial and temporal scales, necessitating innovative approaches to evaluate their performance. Contributions are divided into:

- **Dense Constellations**:
  - **VLEO Constellations**: In [Ott+25], dynamic routing in LEO and Very Low Earth Orbit (VLEO) constellations was analyzed, emphasizing challenges in polar coverage. Supported by Camilla Ottaviani (Politecnico di Torino), in the

context of ESA's HANDING-OVER project, a novel heuristic, *PolarDisCo*, optimized ISL utilization in these regions, reducing end-to-end delays by 14.8%. This work informed broader strategies for routing efficiency and geometric optimization, particularly in Starlink-like configurations. (More in Figure B.2.) VLEO satellites are a promising approach to enhance future mega-constellations. I plan to address these as discussed in Chapter 3.

- **Sparse Constellations**:
  - **Early SARE Model**: The SARE mission, developed during my PhD, exemplifies sparse constellations with episodic communication opportunities. Early studies [Fra+16] analyzed ISL communication and delay-tolerant store-carry-and-forward data handling, proposing the first hierarchical designs for LEO networks leveraging DTN protocols.
  - **Advanced SARE Model**: In the context of the MISSION project, a recent study [Fra+24] evaluated advanced SARE configurations like stair-step and cross-track swarms, identifying trade-offs in memory utilization and ISL ranges under compression and onboard processing. This paper is a project landmark, as highlighted in its website. (More in Figure B.3.)
  - **Sparse RRN with ISL**: I extended the Ring Road Network (RRN) concept, a data mule concept for sparse LEO systems, to reduce latency and improve delivery ratios using in-orbit data transfers. However, challenges like congestion from routing loops were highlighted for the first time, indicating the need for refined congestion control strategies [FFB17b]. (More in Figure B.4.)

2.2.1.2 Constellation Design

Designing satellite constellations involves balancing latency, coverage, and cost constraints across multiple orbital parameters for multiple satellites. Contributions in this domain include:

- **Single-Satellite Missions**: The Maximum Asset Accessibility Algorithm (M3A) [Hen+19], developed by Santiago Henn, my PhD student at CONICET, optimized satellite orbits and ground station placement to maximize access time in the case of Earth observation missions for CONAE. The algorithm achieved near-optimal solutions 23 times faster than an exhaustive search, significantly improving accessibility for LEO relay networks. (More in Figure B.5.)
- **Multi-Satellite Missions**:

- **Global IoT Constellation Design**: A heuristic optimization approach introduced in [Fra+20a] minimized satellite count while maintaining delay-tolerant IoT connectivity unchanged. Results demonstrated global and latency-tolerant LoRaWAN services are achievable with just *nine* satellites, maintaining the 2-hour gap imposed by the protocol definition between contacts.
- **Regional IoT Constellation Design**: Extending global designs, [Cap+22] addressed regional sparse constellations for LoRaWAN and NB-IoT protocols, optimizing configurations for latency and resource use over specific regions (e.g., countries) on the planet. To this end, innovative hexagonal meshing demonstrated superior performance, achieving 3-satellite coverage for Africa while balancing computational demands. (More in Figure B.6.)

2.2.2 Contact Plan Design

This section introduces and advances the Contact Plan Design (CPD) concept, a systematic process for selecting and scheduling communication opportunities in space networks, or *contacts*. **While constellation design provides the physical topology, CPD provides the logical topology, comprising a usable subset of the system contacts**. Defined as temporary links between nodes (e.g., satellites or ground stations), *contacts* are the building blocks of a *contact plan*. A contact plan is a curated schedule of these links optimized to address constraints such as energy budgets, antenna pointing capabilities, interference avoidance, and mission-specific requirements.

Contact Plan Design (CPD). I first introduced the CPD framework in [FF15a] as the process to create the logical contact plan out of a physical contact topology. I grouped key constraints, such as reachable nodes, antenna configurations, and transponder and pointing determinations, which have guided subsequent research. CPD ensures efficient and viable data transfer in resource-constrained space networks.

2.2.2.1 CPD Methodologies

The evolution of CPD methodologies reflects increasing levels of network awareness (see Figure B.11). Most of these schemes were developed during my PhD and extended afterward.

- **Topology-Aware CPD**: Early work in [FMF13] focused on fairness in link assignment, formulating the Fair Contact Plan (FCP) as a multi-objective optimization problem using a weighted matching algorithm based on the Blossom algorithm.

FCP ensures balanced connectivity while maintaining system capacity and scalability, as extended in [FMF14a]. (More in Figure B.9.)

- **Route-Aware CPD**: The Route-Aware CPD (RACP) [FF15b] incorporates end-to-end routing metrics, such as delay, into the planning process. RACP is based on Segmented Architecture (SA) and can accommodate any routing mechanism. By refining FCP through local search techniques, RACP achieves Pareto-optimal solutions, balancing fairness and routing efficiency. (More in Figure B.10.)

- **Traffic-Aware CPD**: The Traffic-Aware CPD (TACP) [FMF16] integrates traffic demand information to minimize delivery delays and optimize resource usage. TACP is implemented as an Mixed-Integer Linear Programming (MILP) model and demonstrated a 42% improvement in resource efficiency and a 10% reduction in delivery times over prior methods. To address TACP's computational complexity, the heuristic TACP-EA[Fra+17c] was introduced, leveraging a custom Genetic Algorithms (GA) formulation for efficient solution generation. (More in Figures B.12 and B.13.)

- **Battery-Aware CPD**: Post-PhD, collaboration with Saarland University extended CPD to include battery dynamics. The Battery-Aware CPD (BACP) [Fra+20b], based on MILP, integrates energy consumption models like linear and kinetic models (i.e., LiBaM and KiBaM) to ensure sustainable network operations. Scalability improvements in [Fra+21] expanded BACP's applicability to larger constellations, addressing challenges in managing energy constraints in nano-satellite constellations. (More in Figures B.14 and B.15.)

- **Position-Aware CPD**: Contributions in [Yan+20] and [NFD22] optimized CPD for Global Navigation Satellite Systems (GNSS), improving key navigation metrics such as Position Dilution of Precision (PDOP) and enabling autonomous distributed CPD approaches for GNSS constellations. Other works in GNSS include evaluations of LEO constellation positioning methods [NDF20].(More in Figures B.16 and B.17.)

2.2.2.2 Tooling: Contact Plan Designer

The Contact Plan Designer (CPD Tool), introduced in [Fra17], bridges academic advancements and real-world implementation. Integrated with Systems Tool Kit (STK), this tool automates CPD by modeling constraints using industrial-grade validate models such as line-of-sight, link budgets, and interference, enabling optimized and applicable contact schedules. The tool's visual and technical capabilities have been critical for engaging

stakeholders, particularly in missions involving SARE, GomX-4, and Skyloom. (More in Figure B.18.)

2.2.3 Mission Operation

The field of satellite mission operations is experiencing a paradigm shift toward automation and optimization. This shift is driven by **the growing size and complexity of modern constellations, which demand scalable, energy-efficient, and autonomous management methods**. My contributions in this domain bridge computer science and space engineering, integrating theoretical advancements with practical applications across diverse areas, including battery-aware operations, advanced mission planning tools, and performance optimization.

2.2.3.1 Automated Battery-Aware Operations

The GomX-4 mission, in collaboration with GomSpace, was a testbed to validate my battery-aware CPD methodologies.

GomX-4 is a dual-satellite cross-linked mission demonstrating automated and energy-efficient network operation. Integrating dynamic programming (DP), antichain pruning, and advanced battery models like KiBaM into the LEOPOWVER toolchain enabled real-time, autonomous scheduling of payload and ISL communication tasks. Key innovations included a receding-horizon approach to dynamically adapt plans based on real-time telemetry, validated over a 48-hour fully autonomous in-orbit experiment [Sto+20]. (More in Figure B.19.)

Key Operations In-Orbit Experiment. Results from the GomX-4 experiment confirmed the feasibility of human-independent satellite operation while optimizing task execution and energy utilization, setting a benchmark for autonomous constellation management for resource-constrained nano-satellites.

2.2.3.2 Advanced Operation Models

Building on the GomX-4 experience, my subsequent research advanced mission planning frameworks, emphasizing polygon-based coverage models, network size estimation techniques, and satellite visibility predictors.

- **Polygon-Based Coverage Models**: This method addressed the limitations of grid-based coverage estimation by employing polygon-based Boolean operations and

transformations between Euclidean and non-Euclidean spaces [HFH23]. The approach models satellite footprints as precise polygons, enabling efficient and accurate coverage calculations, particularly for fine-grained scenarios, surpassing tools like STK in computational efficiency and accuracy. This is one of the core contributions of Santiago Henn's PhD (More in Figure B.20.)

- **Network Size Estimators**: Collaborations with Pablo Illabaca and Diego Maldonado, my PhD student at Inria, produced lightweight estimators such as Optimistic Collision Information (OCI) [Par+23] and its extension for LoRa/LoRaWAN networks (L-OCI) [Mal+23]. These methods optimize DtS-IoT network scalability, accurately and computationally efficiently estimating active IoT nodes in satellite footprints. This output is used by data-driven Medium Access Control (MAC) mechanisms to throttle the channel competition. OCI-based frameworks achieve throughput near theoretical limits and enhance scalability for networks with thousands of IoT devices. (More in Figure B.21.)

- **Satellite Visibility Predictors**: Addressing the energy constraints of DtS-IoT devices, the algorithm in [Ort+24], led by Raydel Ortigueira, my graduated PhD, predicts satellite visibility using a two-body orbital model. It reduces computational overhead by up to 66.53% while maintaining high accuracy, enabling resource-efficient communication for constrained IoT devices. Efficient satellite visibility prediction is crucial in resource-constrained ground devices for DtS-IoT.

- **Parallelizing Orbital Propagators**: My early work [FFM13b] on parallelizing the SGP4 algorithm using OpenCL achieved up to 20x speed-ups, enabling real-time orbital propagation for large constellations. This advancement supports collision avoidance, mission planning for large-scale satellite operations, and onboard propagation execution over emerging multi-core radiation tolerant onboard processors such as LEON-3.

Evolution in the Field of Analysis, Design, and Operations: I introduced innovations like the **Contact Plan Design (CPD) framework**, which optimizes link scheduling under resource constraints, and extended it with battery-aware and traffic-aware methodologies to enhance sustainability and efficiency. My contributions to **constellation design** have optimized sparse and dense networks for global and regional connectivity, providing scalable and cost-effective IoT and Earth observation solutions. I pioneered battery-aware **autonomous scheduling and advanced planning models in mission operations**, validated through real-world experiments such as GomX-4, which set benchmarks for energy-efficient,

human-independent satellite operations. Collectively, these advancements provide a robust network architecture and management framework that bridges theoretical insights and practical implementations.

2.3 Routing and Flow Control

This section focuses on **multi-hop** routing and transport protocols, addressing two key questions: *How to determine the next hop in multi-hop satellite networks?* and *How to manage data flows effectively?*

The contributions are structured into **Low-Earth Orbit Constellations** and **Delay-Tolerant Networks**. This evolution reflects a progression from DTNs, inspired by Argentinian Commission on Space Activities (Space Agency) (CONAE)'s Segmented Architecture for Earth Observation (SARE) mission, to scalable solutions for LEO mega-constellations, combining theoretical innovations with practical applications to address the unique temporal and spatial dynamics of satellite communication systems. This section is further expanded in Appendix C.

2.3.1 Low-Earth Orbit Constellations

As mentioned, LEO constellations are transforming global connectivity (see Figure B.1). Despite their promise, these constellations introduce challenges, including dynamic topologies, frequent link changes, and latency-sensitive requirements. My contributions in this domain focus on advancing **Routing** and **Flow Control** protocols.

2.3.1.1 Routing

Routing in LEO constellations requires innovative approaches to manage dynamic topologies and ensure reliability. My research introduces two key advancements:

- **SATNET-OSPF**: In collaboration with CAST, I developed SATNET-OSPF, an extension of OSPFv3 tailored for integrated space-terrestrial networks [Yan+24a]. Key innovations include a memory-efficient adjacency list, fast link state detection, proactive routing mechanisms, and implementation on space-grade hardware, including SPARC V8 CPUs and radiation-hardened FPGAs. The core value of SATNET-OSPF is that it seamlessly integrates with terrestrial OSPF routing. (More in Figure C.1.)

- **DiscoRoute**: With my PhD student Gregory Stock, I introduced DiscoRoute, a distributed on-demand routing algorithm optimized for mega-constellations [SFH22a; SFH22b]. This method leverages orbital dynamics models to compute efficient paths, minimizing hop count and travel distance with low computational overhead. Being specialized for in-orbit routing, DiscoRoute achieves scalability, being 158× faster than Dijkstra's algorithm and 7× faster than DAG-based methods in Starlink-like mega-constellations. (More in Figure C.2.)

2.3.1.2 Flow Control

The unique characteristics of LEO constellations, such as high end-to-end latency and frequent route changes, necessitate unique transport-layer optimization. My work was one of the first to discuss this need. It focuses on these by co-designing routing and transport layers:

- **Evaluation of BBR in QUIC**: A study led by Yue Wang evaluated Google's QUIC protocol with the BBR congestion control algorithm for satellite Internet [Wan+18]. Testing in emulated Geostationaty Orbit (GEO) satellite conditions revealed that QUIC with BBR outperforms traditional CUBIC algorithms, particularly under high packet loss (>1%) and latency (200–1000 ms). While this work showed that QUIC appealed to satellite networks, it was evaluated on relatively stable GEO systems.

- **Congestion Control and Route Stability**: Addressing the instability caused by frequent route changes in more dynamic LEO constellations, I supervised Gregory Stock in developing routing algorithms such as Stubborn Dijkstra, Tenacious Routing, and SetCover [Sto+24]. These methods leverage predictable orbital dynamics to optimize route stability while balancing end-to-end delay. Tenacious Routing improves TCP data rates by 2× compared to Dijkstra by minimizing disruptive route changes. SetCover minimizes route changes, significantly enhancing stability without substantial performance trade-offs. I deepen on this attractive and impactful routing-transport co-design research line later in the perspectives in Chapter 3. (More in Figure C.3.)

2.3.2 Delay-Tolerant Networks

Delay-Tolerant Networking (DTN)s present unique challenges in routing and transport compared to LEO networks due to their inherent reliance on predictable temporal dynamics. DTNs require data stored onboard during link outages and forwarded when connectivity resumes (see Figure C.4). This is known as store-carry-and-forward data

handling (not to be confused with the store-and-forward term, without the *carry* part, used in layer two switching).

Routing and Transport in DTNs. My work, informed by the foundational insights of Scott Burleigh, underscores the critical importance of **effectively integrating the time dimension into store-carry-and-forward network operations**, a perspective largely absent in the static and persistently connected terrestrial Internet protocols [FFB17a].

2.3.2.1 Routing with Scheduled Contact Plans

Space DTNs benefit from *predictable trajectories* and communication opportunities, enabling precomputed contact plans. My research focuses on Contact Graph Routing (CGR), a framework leveraging this predictability for efficient routing:

- **Contact Graph Routing**: One of my most cited works, the tutorial in [FJB21], systematically addresses routing challenges in time-evolving topologies using contact plans. CGR adapts Dijkstra's and Yen's K-shortest path algorithms to optimize path computation and forwarding. Key innovations in my portfolio include adapting Yen's algorithm for robust path selection. **My solution was integrated into NASA's ION and HDTN flight-grade protocol stacks**, demonstrating CGR's practical utility aboard spacecraft, including the ISS, where these stacks run permanently on scientific payloads. I also proposed advanced route computation methods, such as *one-route* and *per-neighbor*, to address scalability challenges [Fra+18a]. (More in Figures C.5, C.6, and C.7.)

- **Reliability**: My early work revealed CGR's sensitivity to disruptions in complex constellations, driving the development of fault-resilient models [Fra+17b]. These insights established the foundation for enhanced rerouting strategies in distributed DTNs.

- **Applicability**: My studies also demonstrated trade-offs between centralized and decentralized application schemes [Mad+18a; FG21], proposing hybrid strategies for scalable operations in missions ranging from small-scale CubeSat networks to larger interplanetary setups. We also introduced the *Spot of Maximum Knowledge* (SMK) routing strategy to address routing challenges in a novel Lunar RRN concept [FFW18], demonstrating the feasibility of applicable communication networks for future lunar exploration missions.

- **Viability**: In cooperation with Olivier De Jonckere, we proposed a contact segmentation technique, which optimizes volume management by dividing contacts into

time-based segments, enabling more precise bandwidth utilization and improving delivery ratios, delays, and hop counts by up to 16% without added computational overhead [DFB24]. This work is at the frontier of routing and congestion control in space DTNs. (More in Figure C.9.)

- **Scalability**: Scalability is likely one of the most challenging open issues with CGR.
  - **Empirical Scalability Limits**: Empirical and theoretical studies revealed CGR's challenges on resource-constrained hardware, such as nanosatellite onboard computers [VF20], an implementation I led for CONAE. (More in Figure C.10.)
  - **Mitigation via Onboard AI**: To mitigate it, I proposed the GAUSS framework, which employs Graph Neural Networks (GNNs) to accelerate onboard routing, reducing computational overhead by half while maintaining operational performance [Olm+23]. (More in Figure C.11.)
  - **Mitigation via Regionalization**: Another mitigation we led with Pablo Madoery is leveraging regionalization approaches, such as Global Regional Knowledge (GRK) and Local Regional Knowledge (LRK), to partition networks into manageable regions, reducing computational overhead [Mad+18b]. With Olivier De Jonckere, we also developed graph contraction techniques for simplified inter-regional routing [JF23]. (More in Figure C.12.)
- **Shortest-Path Tree Networks**: Shortest-Path Tree Routing for Space Networks (SPSN), a proposal we put forward with Olivier De Jonckere, redefines routing models by replacing contact graphs with *scalable spanning-tree structures*. SPSN addresses one of the primary limitations of CGR: its reliance on generating a contact graph for each source-destination pair. While in its early stages of development, this promising framework is one of the perspectives I highlight in Chapter 3 since it significantly reduces computation times for large-scale DTNs while maintaining pathfinding efficiency [DF20; JFB23a].

Routing in Scheduled DTNs is uniquely challenging due to the complexity of CGR, which acts as both routing and flow control—a unification driven by time dynamics that weaken traditional OSI layer models. As a global leader in CGR research and its standardization in CCSDS' SABR, I embrace the **time dimension as both a critical challenge and a rich avenue for groundbreaking research in space and terrestrial networks**.

2.3.2.2 Routing with Uncertain Contact Plans

Routing under uncertain contact plans extends the deterministic assumptions of CGR to address real-world scenarios where transient faults, interference, or environmental factors disrupt contact predictions. Contributions in this domain include:

- **CGR-2Copies**: Initial work [Mad+18c; Rav+18b], including CGR-2Copies, introduced replication-based strategies that improve delivery ratios and reliability by combining *delivery-time* and *hop-count metrics*. This approach revealed the trade-offs between delivery performance and resource efficiency, setting the foundation for more formal models (More in Figure C.13.)
- **Markov Decision Processes**: Subsequent research leveraged Markov Decision Process (MDP)s to develop theoretical frameworks, such as Best Routing Under Failures (BRUF) and Routing under Uncertain Contact Plans (RUCoP), which optimized routing by modeling uncertain transitions and maximizing delivery probabilities [Rav+18a; Rav+21]. Although RUCoP established a benchmark for reliability, its reliance on global network knowledge posed scalability and applicability challenges, prompting adaptations like L-RUCoP and CGR-UCoP for resource-constrained and local decision-making processes. (More in Figures C.14, and C.15.)
- **Statistical Approaches**: To address scalability and applicability limitations, statistical model checking approaches, including Lightweight Scheduler Sampling (LSS), were explored to enable distributed decision-making based on limited local node information [DFH20; DAr+22]. This is crucial in DTNs, where local observations can hardly include information on the traffic status of remote nodes. LSS-based methods demonstrated comparable reliability to global schedulers while improving scalability.
- **Learning Approaches**: Building on this, learning-based approaches, such as Q-Learning (QL) and its localized variant (L-QL), employed reinforcement learning to balance computational efficiency with delivery performance [DAr+24]. These strategies exhibited strong adaptability and resource scalability in dynamic and complex DTN scenarios, further evolving the field towards robust and practical solutions for routing in uncertain space networks. (More in Figure C.16.)

Routing in predictable topologies is already complex, but **introducing uncertainties significantly amplifies the challenge**, especially if considering multi-copy approaches. Addressing this remains one of my agenda's most demanding short-term research goals, as detailed in Chapter 3.

2.3.2.3 Tooling: DtnSim

DtnSim[4], introduced in [Fra+17a], but extended for uncertain routing afterward, is a discrete-event simulator built on OMNeT++ for evaluating routing algorithms and network behavior in space-terrestrial DTNs. It bridges the gap between flight-grade software (e.g., NASA's ION) and scalable simulation environments, enabling realistic evaluations of large-scale, time-evolving topologies. Key features include integration with actual contact plans, support for NASA's CGR library, and rapid prototyping capabilities, making it a cornerstone tool for advancing DTN routing and congestion research. (More in Figure C.17.)

Recent collaborations [Bul+24] demonstrated DtnSim's role in a comprehensive comparison of emulation platforms, validating protocols and mission scenarios for applications ranging from cislunar networks to the ISS. Its integration with visualization tools like IPN-v (discussed in Section 2.1.2.2) highlights its adaptability and impact on developing standardized emulation frameworks for advancements in cross-disciplinary DTN.

2.3.2.4 Flow Control

Flow control in DTNs demands innovative mechanisms due to high-latency links, temporal storage constraints, and the risk of congestion arising from both contact overbooking and storage limitations. My contributions focus on enhancing **reliable delivery** and **congestion control** in this context, addressing these unique challenges.

Reliable Delivery Collaborations with Prof. Kanglian Zhao and Prof. Ruhai Wang advanced reliable transport-layer protocols like Licklider Transmission Protocol (LTP) to handle long delays and lossy links. Key contributions include:

- **Quantifying File Delivery Time**: We developed models for Bundle Protocol (BP) performance under asymmetric channels [Cao+18], revealing the critical impact of acknowledgment delays on delivery time.
- **Impact of Disruptions**: We analyzed BP and LTP performance under intermittent relay disruptions in International Space Station (ISS) communications, highlighting the need for robust retransmission strategies [Zho+19].
- **Energy-Efficient Delivery**: We explored trade-offs between proactive and reactive retransmission strategies, emphasizing the energy costs of ensuring reliability [Liu+19].

[4]DtnSim public repository: https://gitlab.inria.fr/jfraire/dtnsim

Congestion Control To reduce computational overhead, my early work introduced one of the first congestion-aware routing enhancements like Cache-CGR [FMF14b]. Further contributions addressed:

- **Congestion Analysis**: Early during my PhD, I introduced Cache-CGR (C-CGR), an enhanced CGR-based routing mechanism for predictable DTNs, reducing computational overhead by caching next-hop decisions while maintaining stable throughput in high-traffic scenarios [FMF14b].
- **Congestion Modeling**: I developed new methods such as Multi-Graph CGR (MG-CGR) to manage global congestion through node-specific contact plans, achieving near-optimal resource utilization in predictable LEO networks [Fra+15a]. (More in Figure C.18.)
- **Advanced Congestion Strategies**: Supported by Pablo Madoery, we expanded congestion control mechanisms by proposing Global Path-Aware CGR (GPA-CGR), to preemptively manage traffic [MFF18]. The approach is based on booking contact volumes in advance. The complexity of the solution required evolutionary genetic optimizations, which this paper introduced to make congestion management computationally tractable.
- **Learning-Based Techniques**: The nature of congestion mitigation is ideal for learning techniques. Supported by Prof. Kanglian Zhao's team, we recently introduced reinforcement learning approaches based on Proximal Policy Optimization (Proximal Policy Optimization (PPO)) to adapt to traffic patterns in deep-space environments dynamically, balancing latency, resource usage, and delivery efficiency [Yan+24b]. (More in Figure C.19.)
- **Storage-Congestion Interplay**: In cooperation with Bob Short (NASA), we investigated semiring-based models to optimize storage and communication resources, providing a scalable approach for interplanetary DTNs [Ber+24]. (More in Figure C.20.)

DTN congestion. Like the terrestrial Internet, DTN congestion arises from routing and limited resources along the computed route, requiring senders and receivers to adapt to avoid buffer overflows. However, the **lack of immediate feedback to regulate data injection complicates congestion management significantly**. This presents a promising opportunity for effectively leveraging advanced machine learning techniques to address these challenges.

Evolution in the Field of Routing and Flow Control: I introduced **groundbreaking advancements in Contact Graph Routing (CGR)**, optimizing path computation and forwarding for DTNs and demonstrating their applicability in NASA's ION and HDTN stacks. My work on **scalable routing models** for mega-constellations, such as DiscoRoute and hybrid regionalization techniques for DTNs, has established efficient and resilient frameworks for complex satellite systems. In flow control, I advanced **congestion-aware strategies and machine learning-based solutions**, addressing the lack of immediate feedback in DTNs to mitigate congestion and optimize resource utilization. These contributions have defined new benchmarks for routing and transport protocols, enabling scalable, efficient, and robust space communication systems across LEO constellations and interplanetary networks.

2.4 Links and Medium Access Control

This section addresses the question of *How data transfer and access are efficiently coordinated in space networks?*. It focuses on **point-to-point** communications at the data link layer.

As with the other research axis, my contributions in this domain ensure that space systems' protocols and techniques are grounded in practical and scalable solutions. Moreover, I care about bridging foundational link technologies with higher-level networking abstractions discussed in the previous sections. This section is further expanded in Appendix D.

2.4.1 Links

The section delves into advancements in space-to-space and space-to/from-ground communication technologies, emphasizing the importance of **directional transmission methods for energy-efficient and reliable links**. These foundational contributions are the basis for higher-level abstractions like contact plans and routing protocols.

2.4.1.1 Radio Frequency (RF) Links

- **Proximity-1 Modem**: My work on the STI-PRX-01 modem for CONAE 's SARE mission introduced autonomous and adaptive RF communication capabilities for inter-satellite and Earth-satellite links. The modem implements CCSDS' Proximity-1 Space Link Protocol. Developed in cooperation with Fulgor, the modem leverages

advanced features like LDPC coding and configurable modulation to improve reliability and energy efficiency, setting a new standard for small satellite sparse constellations [FK15]. (More in Figure D.1.)

- **Contact Determination and Tracing**: During the COVID-19 pandemic, I explored close-range RF-based proximity estimation and deep contact tracing using BLE signals, demonstrating how these models could be adapted for space applications [Mad+21; Che+23].

2.4.1.2 Free-Space Optical (FSO) Links

- **Space-to-Ground**:
  - **FSO Availability**: In collaboration with Prof. Kanglian Zhao from Nanjing University, we worked on optimizing FSO network availability under atmospheric disruptions using acMILP and derived heuristic algorithms for relay placement and multi-hop routing, significantly improving operational space-to-ground efficiency [Gon+20].
  - **5G FSO Offloading**: Cooperating with Alexis Dowhuszko, we proposed a hybrid backhaul network combining GEO and LEO satellites with terrestrial FSO technologies to optimize 5G edge content delivery, significantly reducing placing times compared to terrestrial-only solutions [Dow+20]. Results demonstrate the cost-effectiveness and efficiency of small LEO constellations in complementing GEO systems for scalable content caching in geographically distributed regions. (More in Figure D.2.)
- **Space-to-Space**:
  - **FSO Terminal Placement**: In partnership with i2CAT, I advanced terminal placement in the spacecraft body, formulating an optimization problem, which I addressed using SA and GA. This is the first paper in the literature dealing with this relevant problem; I generalized it from real problems faced by Skyloom. These approaches enhance the overall FSO link performance, ensuring enhanced inter-satellite data transfer [FRF24].
  - **Hybrid FSO Evaluation**: Still in collaboration with i2CAT, and led by Elena Fernandez-Niño, a PhD student I co-supervise with Joan Ruiz-De-Azua, we introduced a hybrid ISL communication architecture integrating RF and FSO technologies to quickly reduce pointing uncertainties in dynamic satellite scenarios [Fer+24]. Combining TLE-based predictions with real-time GNSS

data achieves sub-millidegree precision, reducing acquisition time and ensuring reliable high-throughput FSO link establishment.

- **Laser Link Scheduler for Deep-Space**: Led by Jason Gerard (Concordia University), we proposed a Laser Link Scheduler (LLS) that optimizes autonomous scheduling for deep-space FSO communications. Integrating a max-flow capacity model with a dynamic programming-based algorithm on time-expanded graphs achieves efficient contact scheduling via maximal weight matching [Ger+24b].

Space Links. Research on space links must tackle the critical challenge of precise pointing and scheduling over long-range distances, typically within a **point-to-point framework**. This is a lesson I have gleaned from my industrial collaborations, where practical constraints often highlight gaps in academic research. An exception to this is low-data-rate solutions, such as **direct-to-satellite access**, which inherently align with these constraints, as discussed below.

2.4.2 Medium Access Control

This section highlights the challenges of multiple access in space, particularly in the context of **direct-to-satellite access**, where decentralized and reduced-control negotiation strategies replace traditional rigid scheduling due to the lack of continuous communication with remote IoT devices. It presents contributions to **Direct-to-Satellite IoT**, including analyses, enhancements of Low-Power Wide Area Networks (LPWAN) protocols, Non-Terrestrial Networks (NTN) management, IoT-DTN integration, and the extension of multiple access concepts to the cellular domain.

2.4.2.1 Direct-to-Satellite IoT Access

DtS-IoT introduces new complexities in coordinating access across vast coverage areas without relying on centralized, continuous control. This domain integrates standardized LPWAN protocols with satellite-specific adaptations to optimize scalability, energy efficiency, and reliability.

- **Analysis**:
  - **LPWAN Integration**: Recent papers like [LLF24] and [Sca+24] examine the scalability of LoRa/LoRaWAN and NB-IoT in satellite systems. They highlight innovative features like LR-FHSS for collision resistance and the balance between network server placement (onboard satellites vs. ground) to optimize latency and resource utilization.

  - **Machine Learning (ML)**: In [Cho+24], led by Alexander Ylnner, my PhD student at Inria, a taxonomy of ML applications for Satellite IoT is proposed, identifying gaps in integrating specific metrics and addressing satellite resource constraints with methods like Graph Neural Networks and Federated Learning. (More in Figure D.4.)

- **LR-FHSS Enhancements**:

  - **Headerless Frame Recovery**: Advanced modulation schemes like LR-FHSS, recently proposed by Semtech are explored to improve scalability and collision resilience in high-density IoT environments [FGI23a; Mal+24a]. (More in Figures D.5 and D.6.)
  - **Sequence Design and Demodulator Allocation**: LR-FHSS-specific optimization strategies, including Wide-Gap Frequency Hopping Sequences (WGFHS) and early decoding/drop mechanisms, demonstrate significant improvements in network performance [Mal+24b]. This paper was led by Diego Maldonado, my other PhD student at Inria, showcasing the relevance of the STEREO project in my integration. (More in Figure D.7.)

- **Transmission Policies**: Approaches leveraging trajectory-based scheduling [Álv+22] and advanced transmission probability functions (TPFs) [VFH21] enhance scalability and mitigate contention issues, achieving near-optimal throughput in realistic satellite scenarios. These works were developed by Guido Alvarez and Kai Vogelgesang, both master's students I supervised at CONAE and Saarland University, respectively. (More in Figure D.8.)

- **Custom Protocol Design**:

  - Custom protocols like RESS-IoT [Ort+21] and MSDQ [GFC24] address challenges in scalability, fairness, and energy efficiency for dense DtS-IoT deployments. These protocols, developed by Raydel Ortigueira and Jason Gerard, respectively, can capture the specifics of DtS-IoT to outperform the performance of unmodified LPWAN standard protocols. (More in Figure D.9 and D.10.)
  - Frameworks like DRAAF [Ger+24a], also developed by Jason Gerard, a PhD student I co-supervise with Prof. Sandra Céspedes (Concordia University), provides tools to evaluate MAC protocols, integrating realistic orbital mechanics and energy consumption models.

- **Management**: Our paper in [Sin+24] introduces a closed-loop feedback mechanism in a Constellation Management System (CMS) for NTN, dynamically reschedul-

ing tasks based on real-time telemetry like satellite power and memory levels. The modular architecture integrates scheduling and network management, achieving a 10% throughput improvement over open-loop approaches. This work significantly enhances the robustness and adaptability of NTNs for future 6G applications. This is a core topic addressed by Arnau Singla (i2CAT). (More in Figure D.11.)

- **DTN and Access Integration**: The integration of DTN and direct-to-satellite access methods is driven by leveraging the strengths and expertise of both approaches.
  - **DTN of Things**: An early book chapter in [FF18] I developed during my PhD already introduces the Delay-Tolerant Internet of Things (DTN-IoT) paradigm, addressing intermittent connectivity and high-latency environments. It explores integrating DTN principles with IoT to enable scalable and resilient operations across resource-constrained scenarios, overcoming the limitations of continuous connectivity in traditional IoT.
  - **DTN-based MAC**: The paper in [CFR18] examines integrating TDMA-based MAC with Scheduled DTNs for satellite networks. It proposes three cross-layer approaches to optimize slotting and routing efficiency in dynamic LEO constellations, balancing complexity and performance. These approaches highlight the importance of cross-layer information sharing between DTN and the link layer.
  - **CGR in DtS-IoT**: The study in [Mon+24], led by Sebastian Montoya, an inter I supervised at Inria, integrates CGR and DtS-IoT within the FLoRaSat framework to optimize routing in high-latency, resource-constrained environments. Key findings demonstrate that DTN's store-carry-and-forward approach significantly reduces delays, particularly in constellations with ISL, showcasing scalability and efficiency for DtS-IoT systems. (More in Figure D.12.)
  - **LoRa in DTN**: In the context of the MISSION project, my intern student Germán Scapin cooperated in D3TN GmbH to integrate LoRa with the μD3TN protocol stack. As discussed in this blog post, the approach enhanced IoT communication under intermittent connectivity. Two methods—store-and-forward mechanisms and LoRa-specific Convergence Layer Adapter (CLA) for DTN—improve resilience, scalability, and cross-protocol compatibility, paving the way for modular DTN-IoT architectures.

2.4.2.2 Direct-to-Cell Access

The Direct-to-Cell (D2C) domain builds on DtS-IoT principles to establish direct satellite connections with handheld devices.

My initial work in this area, represented by [Pas+24], led by Hannah Pasandi, outlines the technological and operational challenges, such as frequent handovers, high Doppler shifts, and link variability. Key contributions include a taxonomy of D2C advancements in beamforming, spectrum sharing, and dynamic protocol design and insights into industry practices and gaps, particularly in scalability and spectral coexistence strategies. Besides the technical challenges, we highlighted that D2C needs viable business models and effective standardization.

Business Models High deployment costs, particularly in underserved areas, and the unproven potential for mass-market smartphone integration highlight the importance of scalable solutions, such as leveraging Mobile Network Operators (MNO)-Satellite synergies, dynamic spectrum sharing, and flexible pricing models.

Standardization Standardized 3GPP interfaces and global spectrum harmonization are essential for seamless compatibility and collaboration between MNOs, Satellite Network Operators (SNOs), and regulators. **Multi-tenant access** protocols and interoperability frameworks will further enhance scalability, enabling innovation and efficient resource allocation within shared satellite infrastructures.

Evolution in the Field of Links and Medium Access Control: I introduced innovative solutions in link technologies like the **Proximity-1 modem** for adaptive RF communications and **hybrid FSO architectures**, enhancing energy efficiency and reliability in inter-satellite and space-to-ground communications. In MAC, my work on **DtS-IoT established scalable protocols**, leveraging advancements in LR-FHSS and custom designs like RESS-IoT to enable efficient resource use in remote and intermittent connectivity scenarios. Additionally, I pioneered frameworks for **integrating DTN principles with IoT and explored cross-layer designs** to bridge MAC and routing for dynamic LEO constellations. These contributions have laid the groundwork for robust, scalable, and adaptable communication systems for IoT and direct-to-cell applications.

Chapter 3

Perspectives and Vision

This chapter concludes the outlook of this HDR, summarizing the current state and setting the stage for future research directions. The **Perspective** section focuses on *short-term* goals, emphasizing well-defined problems and leveraging the available resources within my team to address them effectively. The **Vision** section adopts a more forward-looking approach, outlining a *longer-term* strategic vision for the research area. It addresses technical advancements and organizational strategies, laying the foundation for sustained innovation and impact.

3.1 Perspectives

This section outlines the short-term goals, building upon the topics discussed, the progress achieved, and my research team's current status and funding perspectives.

3.1.1 On Applications, Architectures, and Environment

The future directions within the Section 2.1 domain can be divided into three main areas: two focusing on applications and one on architectures. As this chapter is centered on application-driven contributions, it naturally presents concrete opportunities for collaboration with industry partners.

- **Space-Underwater Applications** (Corresponding to Section 2.1.1, Appendix A.1.1.2): The interdisciplinary domain of underwater and space communications through the ColdSun project represents a key avenue for my future research endeavors. Building on the foundational work of ColdSun, I aim to expand the collaboration with D3TN GmbH in the context of the MISSION project and grow the scope towards a larger, network-wide vision illustrated and described in Figure 3.1. To achieve this, I am actively pursuing funding opportunities from ESA to facilitate a

second iteration of the project, enabling advanced experimentation and broader system-level insights.

- **Space Cloud Applications** (Corresponding to Section 2.1.1, Appendix A.1.2): The recently funded DarkSol project, also by D3TN GmbH, a key partner in the MISSION project, aligns seamlessly with the methods and frameworks developed in the context of Space Cloud by Camilo Rojas. This perspective investigates the feasibility and design of space cloud computing architectures, leveraging innovative technologies such as DarkSol to address the unique challenges of distributed computing in space by carefully adapting the Kubernetes framework. Plans include exploring opportunities for funding and collaboration through ESA calls, aiming to advance the integration of cloud computing capabilities within space networks.

- **HAPS and IoT Architectures** (Corresponding to Section 2.1.2, Appendix A.2.1.1): Integrating High-Altitude Platform Systems (HAPS) with IoT and satellite systems is also a principal objective of my plans. HAPS presents a unique architectural advantage by enhancing direct-to-access connectivity, owing to their low altitude and enabling fully optical satellite communication by operating above cloud cover. This dual capability, largely overlooked by the research community, offers a transformative space and terrestrial network integration opportunity. Future efforts will include collaborating with Prof. Halim Yanikomeroglu and Pablo Madoery from Carleton University in Canada to explore further and develop these concepts.

3.1.2 On Analysis, Design, and Operations

Future efforts along the Section 2.2 research lines will focus on advancing the analysis of VLEO networks, consolidating and enhancing Contact Plan Design (CPD) methods through comprehensive surveys and potential in-orbit validations, and rigorously testing operational methods for network size estimation and visibility prediction in collaboration with key industry partners.

- **VLEO Analysis** (Corresponding to Section 2.2.1, Appendix B.1.1.1): This perspective deals with the unique analytical challenges of VLEO networks, such as seamless handovers, Radio Resource Management (RRM), and interference mitigation. VLEO constellations are a clear future trend, as evidenced by initiatives like Starlink, OneWeb, and the upcoming European IRIS2 constellation, further motivating this research direction. This effort is currently embedded within the HANDINGOVER project, funded by ESA and conducted in collaboration with Prof. Roberto

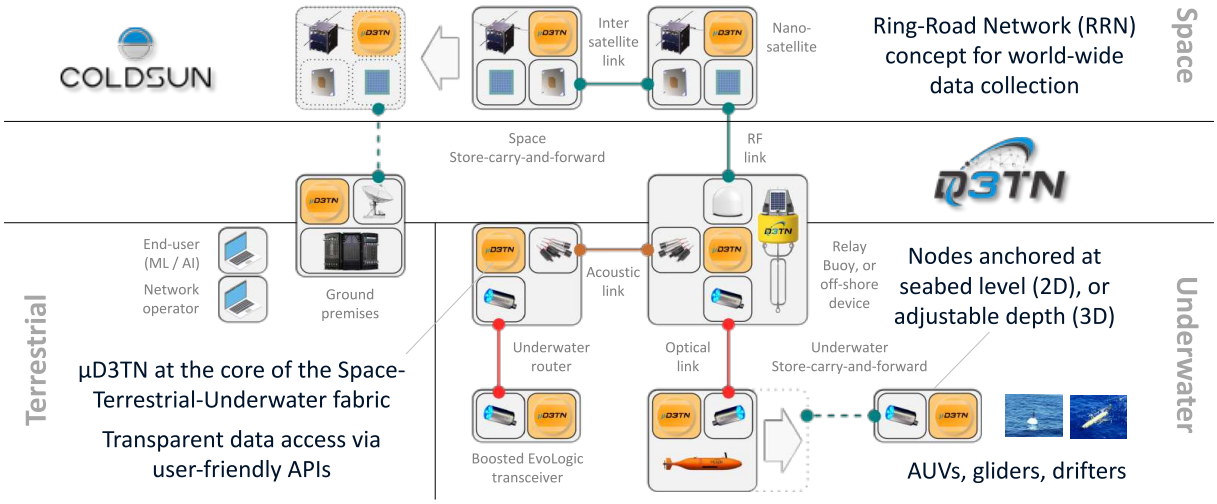

Figure 3.1: ColdSun 2 Concept. Integration of μD3TN into a Space-Terrestrial-Underwater communication fabric for seamless data collection and relay. The Ring-Road Network (RRN) concept utilizes nanosatellites, relay buoys, and underwater nodes to achieve transparent, delay-tolerant data transmission across acoustic, optical, and RF links, supporting diverse applications like AUVs, gliders, and seabed monitoring with user-friendly APIs. (Courtesy from D3TN GmbH)

Garello's team from Politecnico di Torino, providing a robust framework for addressing these challenges.

- **Contact Plan Design Consolidation** (Corresponding to Section 2.2.2, Appendix B.2.2): The goal is to conduct a detailed survey of existing and emerging methods for contact plan design. The survey shall include those developed by other authors and me to identify gaps and opportunities for integrating multiple approaches. Based on the survey, the goal is to develop a comprehensive and standardized CPD library leveraging STK's Contact Plan Designer, refining the framework for broader applicability across missions. Given the maturity of this topic, pursuing an in-orbit validation would be a natural next step, building on my previous successful collaboration with GomSpace, to further solidify the practical relevance and applicability of CPD.

- **Operations In-Orbit Validation** (Corresponding to Section 2.2.3, Appendix B.3.2): This perspective aims to validate methods for network size estimation, visibility prediction, and coverage currently developed by Santiago Henn in operational scenarios for space network design through in-orbit experiments. Throughout my career, I have developed multiple methods in these domains, which can be tested and refined using the infrastructure provided by the STEREO project in collaboration with Kinéis as a key partner. Kinéis 's constellation of satellites equipped with software-defined radios offers a unique opportunity for configurable and accessi-

ble validation.

3.1.3 On Routing and Flow Control

The future directions emerging from Section 2.3 align with advancing transport protocols tailored for mega-constellations and enhancing routing strategies for DTNs, addressing both scheduled contact scenarios and routing complexities under uncertainty.

- **Flow Control** (Corresponding to Section 2.3.1, Appendix C.1.2): This perspective line focuses on advancing routing solutions that improve stability and congestion control, building upon previous work on stable route algorithms. With Gregory Stock, from Saarland University, we are developing Virtual Orbital Transport Emulation (VOTE): a tool to emulate multi-hop routes in mega-constellations, providing a testbed for evaluating diverse congestion control algorithms, including TCP and QUIC. Preliminary results reveal promising optimization opportunities, which we aim to further explore and refine within the framework of ongoing ESA projects with the Saarland team.

- **Scheduled DTN Routing** (Corresponding to Section 2.3.2, Appendix C.2.1): The goal is to tackle the scalability challenges in CGR by advancing regionalization techniques and investigating alternatives like the SPSN. With Olivier De Jonckere from Montpellier University, we are collaborating on a large-scale routing library encompassing various CGR and SPSN variants. This library will enable unprecedented analysis capabilities and facilitate seamless integration into future DTN protocol stacks, including ION, μD3TN, and HDTN, which are currently within the scope of our ongoing research efforts. Another related effort is developing a BP-Socket library to connect DTN applications seamlessly[1].

- **Uncertain DTN Routing** (Corresponding to Section 2.3.2, Appendix C.2.2): This perspective line investigates the challenges associated with partially observable Markov Decision Processes (POMDPs) and their scalability. Currently, I am supervising Alexander Haberl, a Master's student from Saarland University who is developing POMDP models using the JuliaPOMDP framework. The hypothesis is that POMDPs can decide local routing in the presence of dependent contact failures. Most importantly, POMDP models can be used to predict the number of copies in remote areas of the network, thereby supporting the estimation of delivery probability. This work provides a robust foundation for incorporating uncertainty into routing strategies.

[1]BPSocket public repository: `https://github.com/DTN-MTP/bp-socket`

Integration of Scheduled and Uncertain DTN Routing: The ultimate challenge in DTN routing lies in developing a unified approach that seamlessly integrates scheduled and uncertain methods to handle both scenarios simultaneously. While this goal remains distant due to the compounded complexity of the two problems, pursuing this high-risk, high-reward research avenue offers the potential for groundbreaking advancements with practical applications in DTN.

3.1.4 On Links and Medium Access Control

The primary ambitions within the scope of Section 2.4 include advancing hybrid FSO and RF link technologies, consolidating progress in DtS-IoT access with a focus on LoRa, LR-FHSS, and NB-IoT protocols, and delving into the transformative and technically demanding domain of Direct-to-Cell (D2C) access.

- **Hybrid Free-Space Optical Links** (Corresponding to Section 2.4.1, Appendix D.1.2): This perspective focuses on advancing the development of hybrid FSO and RF links, leveraging LPWAN protocols to exchange positioning information for faster pointing and reduced acquisition latencies. Preliminary results, presented by Elena Fernandez-Niño through the i2CAT collaboration, serve as a foundation for further exploration. Future work aims to deepen the analysis and integrate Proximity-1 over the LPWAN-enabled RF side-channel to enhance system reliability and efficiency.

- **Direct-to-Satellite IoT Access** (Corresponding to Section 2.4.1, Appendix D.2.1): The primary objective for the coming year is to consolidate ongoing DtS-IoT advancements led by Diego Maldonado and Alexander Ylnner at Inria. This involves addressing LoRa and LR-FHSS coexistence challenges and proposing custom protocols tailored for Kinéis, intending to achieve in-orbit validation within the scope of the STEREO project, set to conclude in two years. Additionally, this research will extend into the realm of NB-IoT through the i2CAT collaboration, under the PhD co-supervision of Arnau Singla, and in partnership with Sateliot, a company leveraging this alternative but promising DtS-IoT technology. These efforts aim to enhance the applicability and performance of DtS-IoT solutions.

- **Direct-to-Cell Access** (Corresponding to Section 2.4.1, Appendix D.2.2): This ambitious perspective focuses on unraveling the Direct-to-Cell (D2C) paradigm, a transformative development in satellite access primarily driven by industry but marked by a significant gap in public research. The complexity of this domain lies in the unique challenges posed by satellite dynamics and extended ranges, which

push mobile network technologies to their limits. Additionally, I plan to focus on the challenge of splitting Satellite Network Operators (SNO) and Mobile Network Operators (MNO), addressing the complexity introduced by stateful information across base stations and core networks. I aim to build on the foundational work of Hannah Pasandi and actively contribute to this rapidly emerging and impactful research area, addressing the technical intricacies and the lack of publicly available insights.

3.1.5 Future Projects

The following projects are fully defined, meticulously written, and strategically designed regarding resources and methodologies. They are either currently under review or will soon be submitted for consideration in national and international funding calls.

IMPACT The *Interplanetary Models, Protocols, and Algorithms for Connected Things* (IMPACT) project seeks to revolutionize interplanetary IoT communication by integrating Direct-to-Satellite Internet of Things (DtS-IoT) and Delay-Tolerant Networking (DTN) architectures into an interplanetary IoT network (IPN). The project tackles critical challenges in IPN mission design, remote asset autonomy, and cross-layer network protocol architectures. Its objectives include optimizing network configurations using **exact and heuristic methodologies**, enhancing operational autonomy with **machine learning**, and creating **resilient, scalable protocols** for IPN backhaul and distributed IoT operations. IMPACT leverages cutting-edge hardware and software innovations, such as CubeSat platforms, LPWAN adaptations, and advanced simulation frameworks, to enable cost-effective, wide-area planetary access and autonomous interplanetary missions. Through this endeavor, **IMPACT aims to lay the groundwork for the next generation of space exploration networks**.

ASYNCNET The *Asynchronous Networks* (ASYNCNET) project pioneers a transformative model-driven methodology for asynchronous networking, targeting the challenges posed by high-delay and disrupted environments such as space, underwater, and disaster-stricken areas. Integrating delay-tolerant principles with synchronous Internet frameworks leverages advanced **optimization techniques**, **model checking**, and **machine learning** to develop novel **protocols**, scalable **architectures**, and a cutting-edge **emulation testbed**. The project aims to redefine digital connectivity by seamlessly bridging asynchronous and synchronous paradigms. This vision positions **ASYNCNET as a foundational step toward extending the Internet's reach into unconnected domains**, fostering innovation across diverse applications such as space exploration, ma-

rine research, and IoT expansion.

3.2 Vision

This section outlines my research's long-term vision and its broader implications. The **Scientific Vision** defines *what I aim to achieve*—taming the time dimension within networking protocols. The **Applied Vision** outlines *the tangible value and impact*, leveraging cutting-edge informatics to drive advancements in the space industry. Finally, the **Methodological Vision** details *how I plan to achieve these goals*, emphasizing a borderless, collaborative approach and the creation of a globally integrated team that embodies the universal essence of space exploration.

Together, these visions offer a comprehensive roadmap to bridge the **synchronous and asynchronous** networking domains, the **research-to-industry divide**, and the **geopolitical fragmentation** that hinder the creation of an inclusive, scalable, and genuinely universal space Internet. At their core, these three visions aim to bridge and unify these pivotal gaps, each carrying an increasingly complex risk profile, as elaborated in **Vision Risks**. Figure 3.2 organizes the driving questions, goals, gaps to bridge, and risks detailed in the following sections.

3.2.1 Scientific Vision

The work summarized in Chapters A to D converges on developing space networking methods and models that fundamentally address **autonomous decision-making and optimization in the time dimension**. This focus sets it apart from classical network models and methods, which are traditionally static or dynamic but operate with negligible delays and disruptions. From an application perspective, this enables the illusion of stable end-to-end paths, often achieved through advanced techniques like Mobile IP (MIP) and Performance Enhancing Proxies (PEPs). In contrast, such stability is inherently unachievable for sparse or deep-space networks, where delays and disruptions are unavoidable. **This necessitates a novel, time-based modeling approach, which lies at the core of my scientific vision**.

Modeling time in the space networking domain presents a unique challenge. It requires reconciling time's intangible and complex nature with the tangible requirements of data transport across vast interplanetary distances. Unlike terrestrial networks, where time is often abstracted or negligible, **space networking demands explicitly considering time as a dynamic and governing factor** due to space environments' significant delays, disruptions, and relativistic effects. This challenge mirrors humanity's broader struggle to define and understand time, drawing parallels from Newton's absolute time,

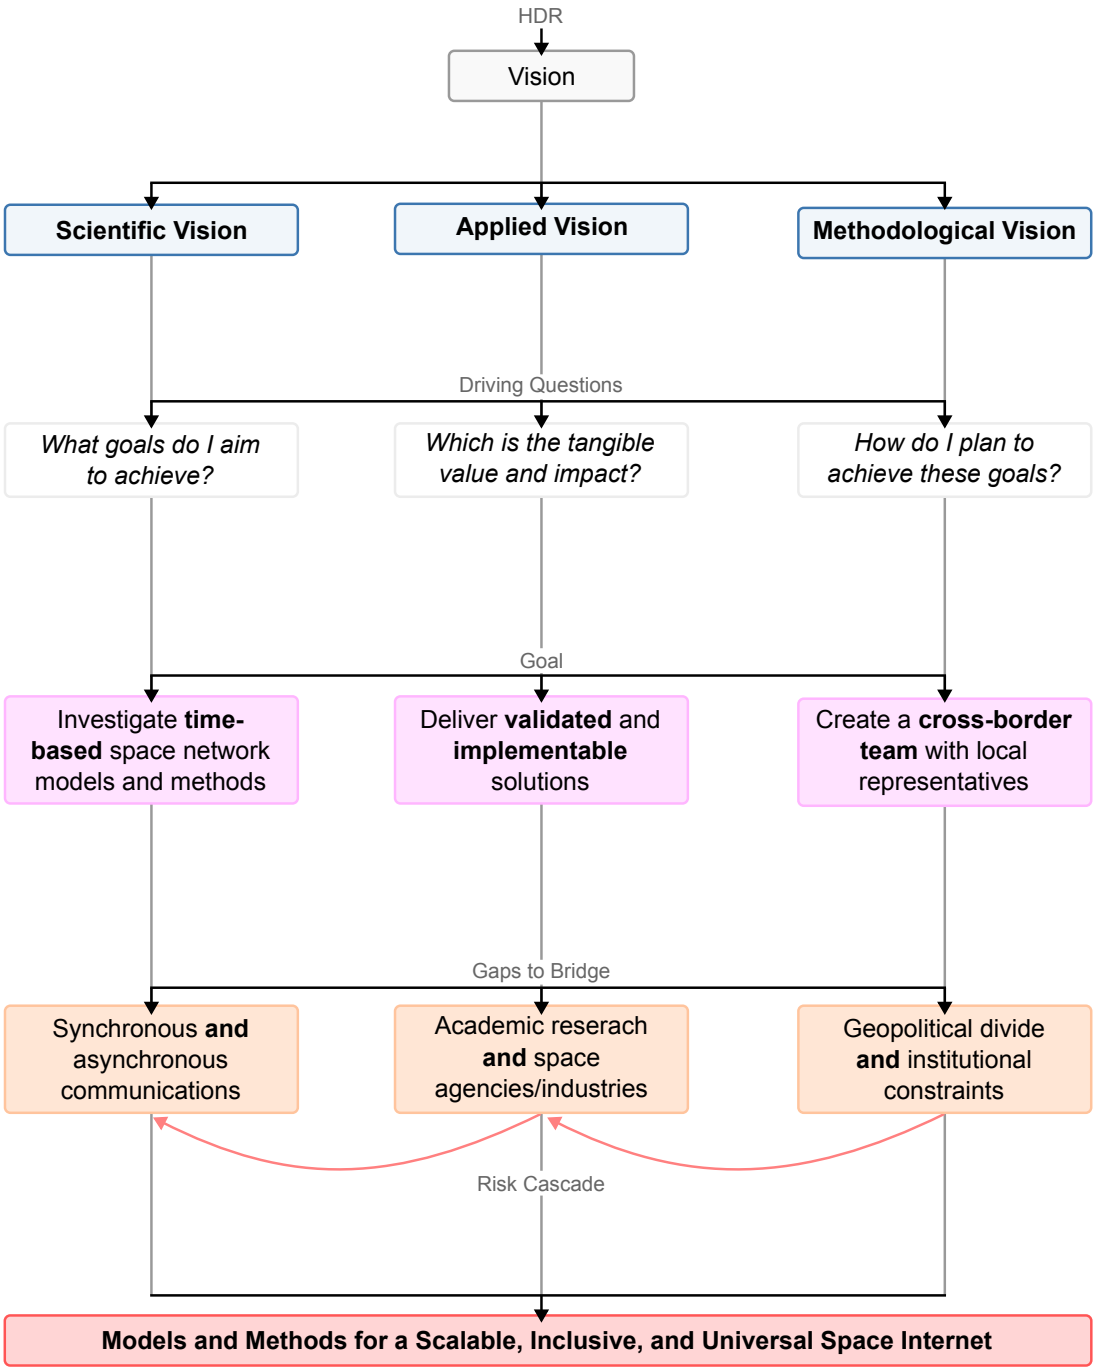

Figure 3.2: Overview of the HDR Vision: The interconnected Scientific, Applied, and Methodological Visions addressing key gaps in space networking, their corresponding goals, and the cascading risks across these domains.

Einstein's relativistic time-space continuum, and contemporary theories that question the very existence of time as a fundamental entity.

The divide between terrestrial and space networks can be framed as the **distinction between synchronous and asynchronous communications: the former assumes immediacy in the response from the next hop or end node, while the latter does not**. The Internet, deeply rooted in the synchronous paradigm, inherits this characteristic from its telephony-based origins, designed for real-time voice communication. Paradoxically, even asynchronous applications like email often operate over inherently synchronous protocols such as TCP, leveraging the abundant synchronous resources available in terrestrial networks. This contrast between synchronous and asynchronous paradigms highlights fundamental, often implicit, assumptions embedded in Internet protocol design. With their unique constraints, **space networks force us to confront and rethink implicit assumptions**, revealing the philosophical underpinnings of how we conceive communication systems and pushing the boundaries of our understanding of networking.

Asynchronous to Synchronous and back to Asynchronous: Early episodic asynchronous communication methods, such as letters, laid the groundwork for establishing the first synchronous connections in telephony. These connections later evolved into the Internet, an outgrowth of telephony networks. Interestingly, as we extend our reach into deep space, the evolution takes an inverse trajectory. We transition from the persistently connected, synchronous terrestrial Internet to an asynchronous space Internet, constrained by the physical realities of planetary occlusions and the finite speed of light over vast interplanetary distances.

Scientific Vision: Space networks explicitly demand a paradigm shift to account for time as a dynamic, governing factor. My scientific vision centers on developing space networking models and methods that fundamentally integrate **decision-making and optimization within the time dimension**, addressing the unique challenges of delays and disruptions inherent to space environments. This goal compels us to rethink the synchronous Internet assumptions and redefine the foundations of a broader network vision, including asynchronous communications.

3.2.2 Applied Vision

I believe informatics is the key to unlocking transformative advancements in space exploration and applications.

As the space sector evolves, the intersection of computer science and space sys-

tems is becoming increasingly pivotal, particularly in areas where research lags behind industrial progress.

- In the **near-Earth** domain, SpaceX's Starlink and competitors like OneWeb are actively advancing in areas such as operations, routing, flow control, inter-satellite links, and direct-to-cell satellite access—several topics discussed in this document. However, their methodologies and technologies remain proprietary, limiting broader insights or collaboration.
- Similarly, in the **deep-space** segment, media speculation about a potential Marslink solution for networking in Mars networks hints at SpaceX's interest in addressing the challenges of interplanetary communication. Despite the evident need for such systems, to my knowledge, there has been no engagement with the scientific community to explore or validate these approaches.

I hypothesize that the adaptable and increasingly open culture within computer science—reflected in the proliferation of open-access research software, methods, and tools—presents a promising path to bridging this gap. **Innovative, performant, and implementable computer science models and methods could effectively promote transparency and induce collaboration with private and government industries**.

> **Warning About the AI Trend**: I caution the research community about the growing trend of publishing space-related AI papers that are often impractical or unexplainable in real-world applications. While AI is undeniably transformative and heavily funded, **the current wave of "AI for space" research risks prioritizing trends over substance**, leading to unvalidated or speculative work-filling conferences and journal slots. By its very nature, space networking demands a more rigorous and conscious approach to AI application—one that prioritizes explainability, reliability, and real-world feasibility over fleeting trends.

This vision aligns with and is empowered by the broader trend toward integrating artificial intelligence and data-driven approaches in space technologies. However, the **focus must remain on precision and explainability, avoiding the pitfalls of opaque or purely probabilistic models**, which are unsuitable for the high-stakes context of space missions. The successful democratization of computational frameworks can accelerate the adoption of informatics in the traditionally conservative space industry, making cutting-edge methodologies more accessible and scalable.

As a result, the long-term priority is to harmonize and empower research and industrial efforts, **ensuring that advancements in space informatics are based on solid theoretical foundations and are practically validated through practical implementations**.

This vision aims to create a new era of innovation by leveraging computer science's collaborative spirit. It will drive the development of scalable, reliable, and efficient space networks that serve exploration and Earth-bound applications.

Applied Vision: By advancing scalable, precise, and explainable AI-driven methodologies, this vision aims to address the unique challenges of space missions while **accelerating the adoption of informatics in the traditionally conservative space industry**. This approach seeks to democratize access to advanced technologies, driving exploration and terrestrial application breakthroughs.

3.2.3 Methodological Vision

Unlike many other research fields, where localized teams and regional collaboration suffice, I believe space networking demands a fundamentally different methodological approach. Its inherently global nature requires a framework that transcends borders and integrates expertise from various nations and continents. **I envision establishing an international research team where local representatives work closely with their respective space agencies, industries, and academic institutions** to foster a shared vision of space networking. This team would bridge geographical divides, promoting consensus and collaboration reflecting space's universal nature.

The terrestrial Internet offers a valuable analogy. In its early years, particularly during the 1980s and early 1990s, the Internet evolved independently across regions due to the absence of a unified communication protocol. The Protocol Wars exemplify this fragmentation, as competing standards vied for dominance until the gradual adoption of TCP/IP as the global standard brought convergence. **Similar "wars" could be avoided in emerging domains like space networking**, which, with its unique global nature, is well-positioned to bypass fragmented growth, making collaboration and consensus essential from the outset. The Internet provides a ready-made infrastructure for real-time global collaboration and a dynamic testbed for live experimentation and development. Moreover, the COVID-19 pandemic demonstrated that many activities, particularly those rooted in informatics, can be effectively conducted remotely. These lessons highlight the potential for **international teams to collaborate across time zones and physical locations**, building frameworks that represent the collective input and interests of a diverse, globally distributed research community.

With these insights in mind, **my goal is to create a cross-border and cross-institutional team that reflects the core principles of the space Internet itself**: inclusivity, connectivity, and universality. This team will pool resources, knowledge, and perspectives from across the planet to develop a space Internet that mirrors the principles of the terres-

trial Internet but extends its scope to near-Earth and deep space. This organizational approach will accelerate innovation and ensure the technologies and protocols align with diverse needs and ambitions. Doing so will lay the foundation for a space Internet based on a frontierless infrastructure that will benefit everyone.

It is worth emphasizing that this approach could be strategically pivotal for Europe's position in the global context. Amid the pseudo-space race between the US and China, Europe risks being relegated to a regulatory role—a perception increasingly echoed in public discourse and memes alike. **A bold initiative to position Europe outside of this gap but at the forefront of space networking research and innovation would counter this narrative and reaffirm its relevance**. This aligns with ESA's vision for the Solar System Internet (SSI), exemplified by the recently initiated Pre-Phase A SSI project, in which I am actively involved. Europe's intellectual and organizational capabilities uniquely position it to take a leading role in shaping and defining this transformative field.

Methodological Vision: Leveraging the Internet's ability to foster real-time global collaboration and lessons from the pandemic on remote productivity, I aim to **build a cross-border research team** that embodies the principles of the space Internet itself—connectivity, inclusivity, and universality—ensuring its development aligns with the diverse ambitions of a globally connected community.

3.2.4 Vision Risks

The risks I foresee in pursuing this vision are not predominantly technical. While technical challenges are complex, they can be addressed with existing informatics frameworks and tools, which are flexible and adaptable to evolving space networking problems. The technical landscape is robust, and the problems are well-defined, making them solvable through diligent research and development. The main hurdle lies within the methodological and organizational aspects of my vision.

From an institutional perspective, **assembling and sustaining a cross-border team introduces significant challenges**. Funding originates from diverse sources with distinct conditions, priorities, and timelines. Institutions often operate under varying bureaucratic constraints that can limit agility and innovation. The growing complexity of geopolitics poses significant challenges to fostering international collaboration, particularly with China—a pivotal partner in the space domain. While Inria provides me with a high degree of flexibility, I remain uncertain if this supportive environment will be sufficient to navigate the obstacles inherent in building such an ambitious and inclusive framework.

Interestingly, these challenges echo the risks we discuss at IPNSIG and IGF concern-

ing the governance of the space Internet. The most significant hurdles are human and institutional rather than technical. Policy, cooperation, and alignment issues across nations and organizations present formidable barriers. **The most complex problems are not in coding protocols or designing architectures but in achieving consensus and overcoming organizational complexities**.

Consequently, the risks inherent in my methodological vision reverberate into the applied vision, where the gap between academia and industry is widening, exacerbated by the current geopolitical climate. Furthermore, these risks cascade into the scientific vision outlined earlier. **This risk potentially triggers an unintended chain reaction, cascading from the methodological vision to the applied vision and, ultimately, undermining the scientific vision**.

Without consensus, the dream of a truly connected space Internet could fracture, leaving us isolated, underperforming, and vulnerable—unable to extend humanity's reach beyond near-Earth orbit.

Success in mitigating or avoiding these risks altogether may be the essential precursor to realizing the inclusive, global space Internet I envision. Addressing these human and institutional challenges could lay the groundwork for a transformative shift in our approach to space networking. We could move from fragmented efforts to a cohesive, borderless framework that mirrors the nature of space itself.

Vision Risks: **These human and institutional challenges mirror the governance difficulties of the space Internet itself**, underscoring that the greatest obstacles lie in fostering collaboration and achieving consensus.

If I succeed in dodging these risks, it's going to be one heck of a ride—to the Solar System Internet and beyond! 🚀

Appendices

Appendix A

Applications, Architectures, and Environment

In this section, I present my research contributions that tackle the core questions driving space exploration and networking: *Why use space?*, *How can space be utilized effectively?*, and *What are the characteristics of the space environment?*

My work in this section spans applications that emphasize the transformative potential of space-based services and architectures that redefine how networks function under the unique constraints of space environments. These contributions systematically challenge and adapt existing paradigms, ranging from extending connectivity to remote regions on Earth through space-based data transport to enabling robust interplanetary and interstellar networks. Through rigorous simulation, experimentation, and architectural innovation, I have laid the foundation for novel *Applications* and *Architectures* that push the boundaries of terrestrial and extraterrestrial communication systems.

A.1 Applications

Applications address the fundamental research question: *why use space?* This compelling inquiry requires creative approaches, rigorous analysis, simulation, and experimental validation to explore and justify its answers.

Contributions span two key application domains: **Space Network Services**, which focus on enabling data transport across satellite and space infrastructures to connect remote devices to the Internet, and **Space Cloud Services**, where processing servers within the space infrastructure handle user demands locally, eliminating the need for round-trip communication with terrestrial networks and significantly reducing latency.

A.1.1 Space Network Services

Space Network Services are a cornerstone of modern satellite communication. They are designed to enable seamless data transport across space-based infrastructures. These services eliminate reliance on terrestrial wired networks, providing connectivity to remote and underserved areas while extending coverage across and beyond the Earth's surface.

> **Space Network Services** facilitate data transport via space-based infrastructures, eliminating reliance on terrestrial wired networks. This approach extends coverage to underserved and remote regions, including oceans and the poles.

In this context, my contributions analyze and collectively enhance the reliability and performance of satellite-based data transport systems, enabling new applications from Earth **observation data transport** to **user data transport**.

A.1.1.1 Observation Data Transport

My contributions to Earth Observation Data Transport address both optical and radar imaging.

Optical Data Handling via DTN In collaboration with Pablo Madoery, the early work in [MFF16] focuses on optical spectrum imaging, proposing segmented satellite constellations to enhance revisit times, coverage, and cost efficiency while tackling the challenges of DTN

S-AIS for the PANSAT Mission The paper in [Alv+21], led by Guido Alvarez, my PhD student, conducts a comprehensive trade-off analysis of Satellite Automatic Identification System (S-AIS) in the context of the PANSAT CubeSat mission, evaluating key aspects such as antenna type, field of view, attitude control, and data processing strategies. The region of interest modeled in Systems Tool Kit (STK) is presented in Figure A.1. This study was used as input for designing the PANSAT mission, which is currently under development by the government of Panama.

SARE Analysis in MISSION In contrast, [Fra+24], recently developed under my lead within the MISSION project. The paper addresses the problem of quantitatively assessing SA for Earth observation concepts proposed by CONAE. This mission aims to monitor maritime zones near Argentina's Exclusive Economic Zone (EEZ) for illegal

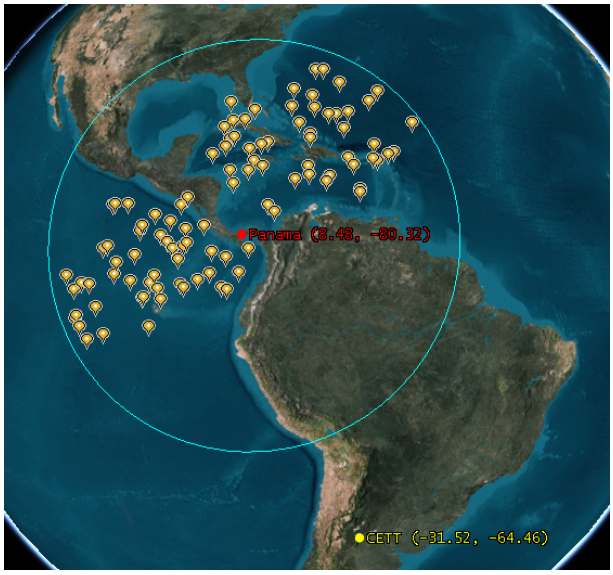

Figure A.1: Region of Interest of the PANSAT CubeSat mission. Yellow markers indicate the position of vessels equipped with S-AIS transponders entering and exiting the Panama Channel. On the bottom, the Teófilo Tabanera Space Center ground station can be observed. (From [Alv+21])

fishing activities, leveraging small, cooperative satellites equipped with Synthetic Aperture Radar (SAR) payloads. To address this, we developed a comprehensive simulation framework using STK, which integrates models for ground and space segments, orbital mechanics, inter-satellite communication, and operational protocols. Various swarm configurations (e.g., along-track, cross-track, stair-step) were analyzed to evaluate coverage, revisit times, contact durations, and data handling efficiency. Advanced modeling techniques included detailed SAR payload parameters and communication link budgets. Key outcomes revealed that configurations with larger angular separations (e.g., 8 degrees) provided more uniform coverage and shorter revisit times, making them suitable for consistent maritime monitoring. Figure A.2 illustrates the analyzed topologies in STK, while Figure A.3 plots the observed contact dynamics. The constrained data handling model also demonstrated the effectiveness of DTN protocols and on-board processing to reduce data volume and optimize transmission efficiency.

A.1.1.2 User Data Transport

My contributions to End-User Data Transport services are encapsulated in two innovative service models centered around the RRN and MCSS concepts.

Ring Road Networks The paper in [Fel+22], also developed in the framework of MISSION, demonstrates the practical feasibility of DTN for User Data Transport using Euro-

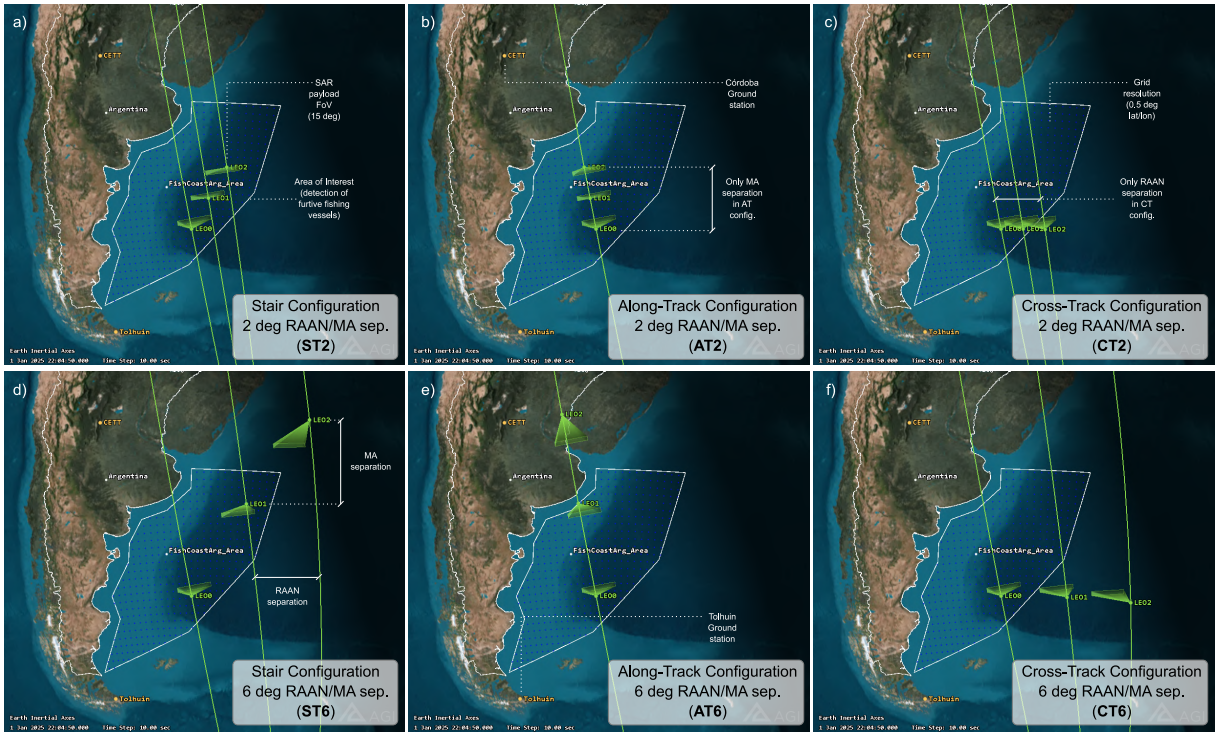

Figure A.2: SAR Mission Models including Stair (ST), along-track (AT), and cross-track (CT) configurations for two and six-degree RAAN/MA separation for the space segment. Ground stations (Teófilo Tabanera Space Center and Tolhuin) and areas of interest (Fish Area) are illustrated in all subfigures with relevant visual information. (From [Fra+24])

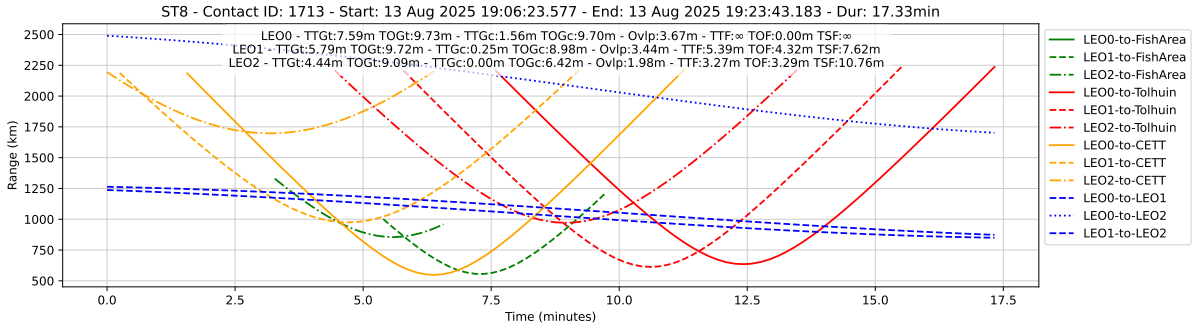

Figure A.3: Contact range dynamics for the ST8 scenario configuration are illustrated in these plots. The green curves represent the distance to the centroid of the fishing area, which serves as the primary region of interest. The orange and red curves show the range to the ground stations. Blue curves depict inter-satellite ranges, with distances between immediate neighbors falling between 700 and 1300 km and distances between the fleet's extremes ranging from 1700 to 2500 km. (From [Fra+24])

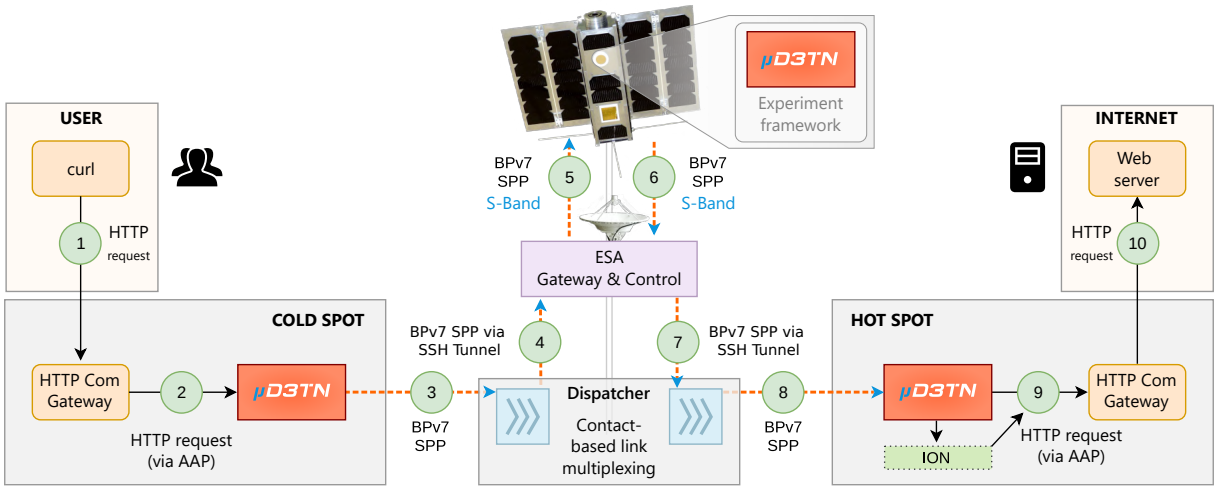

Figure A.4: OPS-SAT experiment setup. The *cold spots* represent isolated nodes requesting Internet services. These nodes utilize the μD3TN DTN implementation to encapsulate HTTP requests within Bundle Protocol version 7 (BPv7), which are forwarded to ESA's OPS-SAT. Equipped with μD3TN on board, the satellite autonomously routes the requests to and from ESA's ground gateway, connecting to the Internet via a designated *hot spot*. (From [Fel+22])

pean Space Agency (ESA)'s OPS-SAT mission as a testbed. This work was developed at an early stage of the MISSION project, funded by the European Union (EU). Through in-orbit experiments illustrated in Figure A.4, it validates the μD3TN protocol, showcasing its capability to enable asynchronous connectivity in remote or underserved regions, thereby advancing the RRN concept with a focus on scalability and real-world applications.

In the context of the **OPS-SAT mission**, I actively contributed with D3TN GmbH to in-orbit experiments testing μD3TN, a lightweight implementation of the DTN protocol. The first experiment on December 11, 2020, successfully conducted the inaugural in-flight test of the BP Version 7, enabling the seamless exchange of Bundles between a μD3TN instance on Earth and another onboard OPS-SAT via ESA's ESOC infrastructure. This milestone also marked the first-ever duplex communication with OPS-SAT over the Internet. Subsequent experiments expanded the scope to include advanced HTTP "ping-pong" message exchanges for delay-tolerant web browsing and validated interoperability between μD3TN and NASA's ION DTN implementation. Furthermore, integrating Google Cloud Vision services allowed the classification of images transmitted from remote cold spots through the OPS-SAT satellite, showcasing the practical and transformative potential of DTN-enabled services for space-based applications.

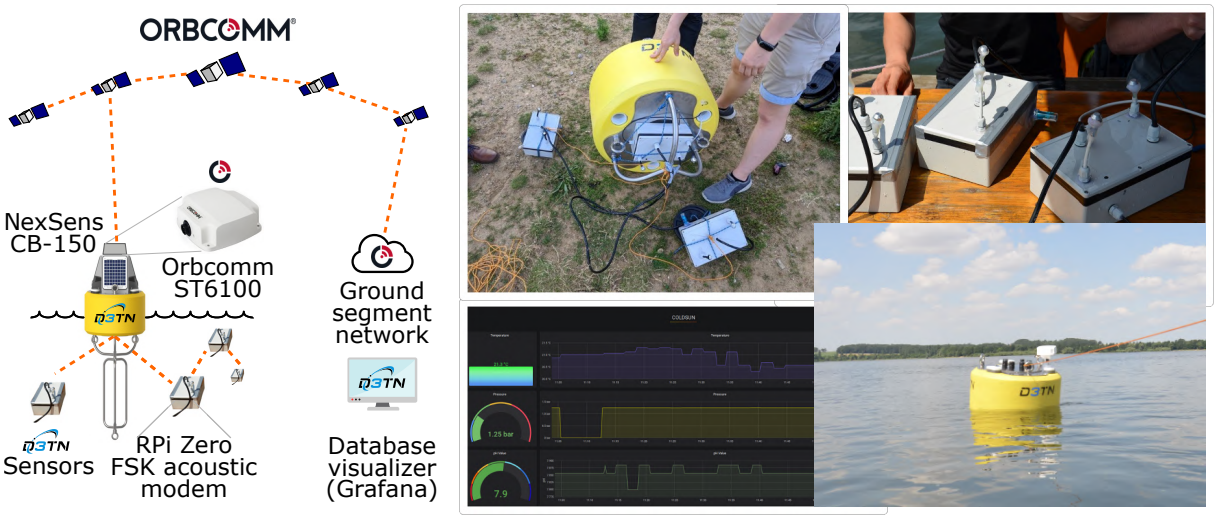

Figure A.5: ColdSun Experiment Overview. Demonstrates seamless data transmission from underwater sensor nodes to the Internet via an acoustic link to a buoy and satellite communication using the μD3TN bundle protocol. The field test validated stable performance in dynamic conditions, showcasing a cost-effective solution for underwater-to-satellite data transfer.

Underwater to Space Communications The ColdSun experiment, in which I also participated with D3TN GmbH in 2019 during my period at CONICET, demonstrated the feasibility of a low-cost and open system for transmitting environmental data from underwater networks to the Internet via satellite communication. The field test at Talsperre Bautzen in Germany involved deploying a buoy equipped with an acoustic communication interface and an Inmarsat satellite terminal. Sensor nodes placed at 1 and 2 meters depths collected pH, temperature, and pressure data, which were transmitted acoustically to the buoy. Figure A.5 illustrates the setup. The buoy bundled the data using the μD3TN BP implementation, buffering and forwarding it via the Inmarsat satellite network to a remote monitoring system. Despite environmental challenges such as movement at up to 2.3 knots and acoustic noise from a boat engine, the system achieved stable and continuous data transmission over two hours. While not reported in a scientific paper (only in ESA's ColdSun website and D3TN GmbH blog posts), this experiment validated the technical solution's robustness, paving the way for scalable and cost-efficient underwater-to-satellite communication systems. It also highlighted the system's potential for maritime industries, including research institutions, offshore operators, and environmental monitoring organizations.

Cross-Agency Communications In 2022, I participated in possibly one of the most ambitious experiments in my career: the Lunar IceCube DTN interoperability experiment. The Lunar IceCube CubeSat, developed by Morehead State University as part of NASA's Artemis-1 mission, had the mission objective to orbit the Moon at low alti-

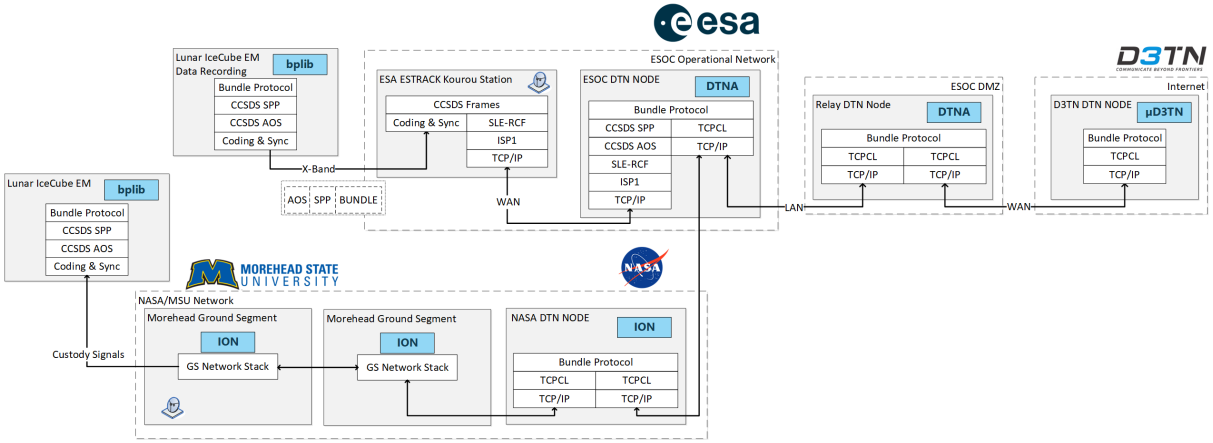

Figure A.6: Cross-Agency DTN Interoperability Setup for Lunar IceCube Simulation. This diagram illustrates the interoperability testing setup for the Lunar IceCube CubeSat mission, leveraging multiple BP implementations across NASA, ESA, and D3TN nodes. Pre-recorded data from Lunar IceCube was replayed via a spacecraft simulator at ESA's Kourou ground station. The data flow began with CCSDS frames received at Kourou, decoded, and forwarded to ESOC via Space Link Extension (SLE) services. ESOC extracted the BP bundles and relayed them over TCPCL to NASA's ION DTN nodes. Extracted data was returned to ESA and forwarded to D3TN's node through a relay node in ESOC's DMZ. (From D3TN blog post)

tudes to investigate the presence of water, ice, and other resources on the lunar surface. However, after launch, communication issues prevented direct communication with the CubeSat. Instead, we emulated the mission for testing purposes using pre-recorded spacecraft data replayed through a generic spacecraft simulator at ESA's Kourou ground station. This emulation enabled critical interoperability tests involving multiple DTN protocol implementations across JPL (NASA), ESA, and D3TN GmbH nodes, showcasing the potential of DTN for enabling robust communication in future space missions.

The **Lunar IceCube Cross-Agency Experiment** was a collaborative experience with JPL (NASA), ESA, and D3TN GmbH. It validated the interoperability of multiple BP implementations in a simulated space environment using the Lunar IceCube CubeSat mission. The setup, illustrated and described in Figure A.6, demonstrated seamless communication across four BP implementations: NASA's bplib on Lunar IceCube, ESA's Java-based DTN nodes, NASA's ION implementation at JPL and Morehead State University, and D3TN GmbH 's μD3TN.

Mega-Constellation Services in Space The paper in [Cap+24] was led by Gabriel Maiolini Capez, a PhD candidate I co-supervise at the Politecnico di Torino. This landmark work

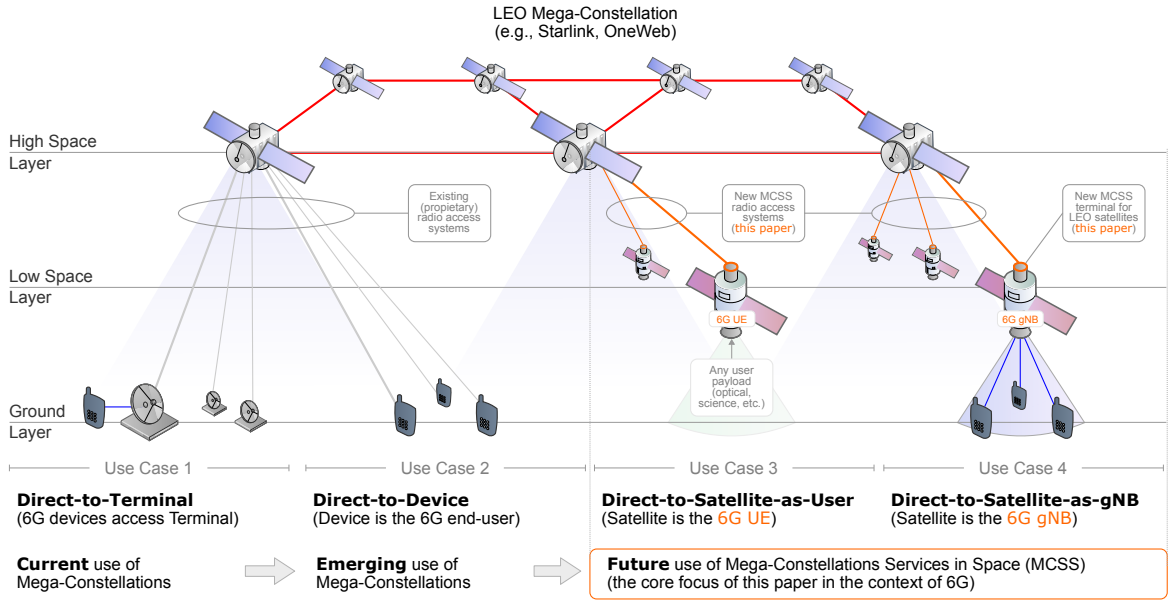

Figure A.7: Mega-Constellation Services in Space (MCSS) paradigm in the context of 6G NTN Networks. Illustrated use cases include 1) direct-to-terminal (e.g., Starlink), 2) direct-to-device (to a smartphone), 3) satellite-as-user (satellite is the end device), and 4) satellite-as-gNB (satellite is a base station serving ground users). The MCSS analysis in [Cap+24] covers the latter two use cases. (From [Cap+24])

introduces the MCSS paradigm, which integrates LEO satellites into 6G NTN, enabling continuous, high-rate communication for space assets. Figure A.7 depicts the proposed service architecture, where LEO satellites leverage higher-orbit mega-constellations for efficient data offloading. Using Monte Carlo simulations and a novel Space User Terminal (SUT) design, the study demonstrates that MCSS outperforms traditional Ground Station Networks and Data Relay Systems, achieving Terabit-scale daily data capacity with low latency and power-efficient architectures.

Direct-to-Satellite IoT Another key aspect of user data transport is the emerging direct-to-device access paradigm, encompassing both IoT connectivity (e.g., Direct-to-Satellite IoT or DtS-IoT) and broadband services (e.g., Direct-to-Cell or Direct-to-Smartphone, DTC or DTS). While not an application *per-se*, notable contributions have been made in these areas [FIV22; FCA19; Fra+22a; Fra+22b], warranting a dedicated subsection for architectures in Section A.2 and technologies in Section D.

A.1.2 Space Cloud Services

This work represents the core research of my current PhD student, Camilo Rojas, co-supervised by Fabio Patrone and Prof. Mario Marchese from the University of Gen-

ova. Our contributions in this domain advance the orchestration and emulation of NTN through innovative Space Cloud Service frameworks and algorithms, leveraging Multi-access Edge Computing (MEC) and Software-Defined Networking (SDN).

Space Cloud Services enable computation tasks to be resolved directly on satellites acting as servers without relying on the terrestrial internet. This approach significantly reduces latency, benefiting delay-sensitive applications while incurring additional costs related to deploying and operating on-orbit servers.

Key contributions in this area include **server orchestration**, where optimal on/off states for edge servers are determined, and **server placement**, which focuses on selecting specific satellites within the constellation to host these servers effectively. These advancements collectively improve the trade-offs between performance and cost in space cloud environments.

A.1.2.1 Server Orchestration

We developed Fuzzy Logic (FL) and Reinforcement Learning (RL)-based orchestration mechanisms for autonomous MEC server activation in LEO constellations [Roj+24b]. The proposed Machine Learning (ML) controllers minimize edge server active time while maintaining low task failure rates, crucial in energy-constrained NTN satellite platforms. The Pareto analysis shown in Figure A.8 evaluates key performance indicators, including server active time and task failure rate. The study demonstrates that FL and RL controllers achieve a superior balance between resource efficiency and service quality under varying task loads. This surpasses baseline approaches such as *Access*, which activates the nearest accessible satellite for every incoming task, and *Bernoulli*. This stochastic method activates edge servers based on a fixed probability. These results proved there is room for improving the performance-cost trade-off in Space Clouds.

A.1.2.2 Server Placement

We measured the latency trade-offs between terrestrial and space cloud architectures across varying cloud server placement strategies [Roj+25]. Figure A.9 illustrates the average network latency ($\bar{\lambda}_\tau$) as a function of the geodesic distance between the devices demanding the cloud service and the fixed ground gateway. Terrestrial clouds exhibit an apparent increase in latency with distance, as computation requests must traverse the satellite network to reach servers on the Internet and return. In contrast, space cloud

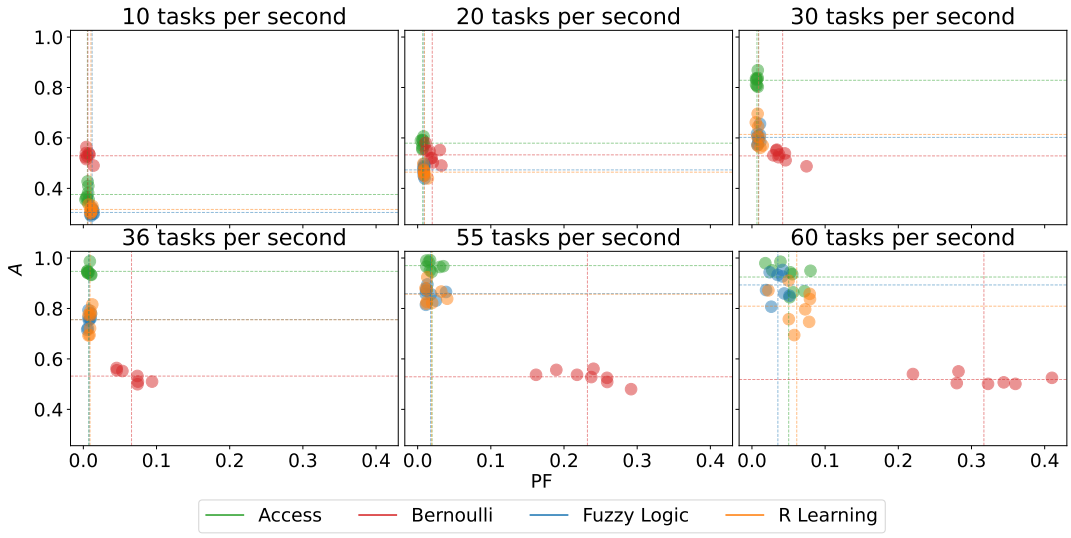

Figure A.8: Comparison of task failure percentage (PF) and overall MEC server active time (A) for Fuzzy Logic (FL), Reinforcement Learning (RL), *Access*, and *Bernoulli* controllers under varying task load configurations. The lower left in each scatter plot, the better. (From [Roj+24b])

configurations, with servers hosted onboard selected satellites within the constellation, provide relatively stable performance regardless of the device's location. Notably, space clouds outperform terrestrial clouds beyond 5000 km. Figure A.10 shows the equatorial plane server distribution configuration. The orbital plane server distributions assume servers are placed in a single orbital plan. These findings highlight the potential of space clouds in reducing latency for globally distributed edge computing tasks and underscore the importance of server placement strategies in NTNs. We currently have a journal paper under review on the latest research in this direction.

A.1.2.3 Tooling: MeteorNet

We developed MeteorNet[1], an emulation tool designed explicitly for continuous-time scenarios to achieve the abovementioned results. MeteorNet's foundational architecture, described in [Roj+23], integrates realistic orbital dynamics with containerized computation environments, providing a robust platform for experimentation and validation of real edge computing strategies in NTN.

Building on this framework, MeteorNet is positioned to generate synthetic datasets [Roj+24a], addressing the critical lack of real data in space NTN. These datasets are instrumental in developing machine learning-driven algorithms for resource allocation and task offloading, bridging the gap between theoretical frameworks and practical applications.

[1]MeteorNet public repository: https://gitlab.com/camilo.rojas/meteornet

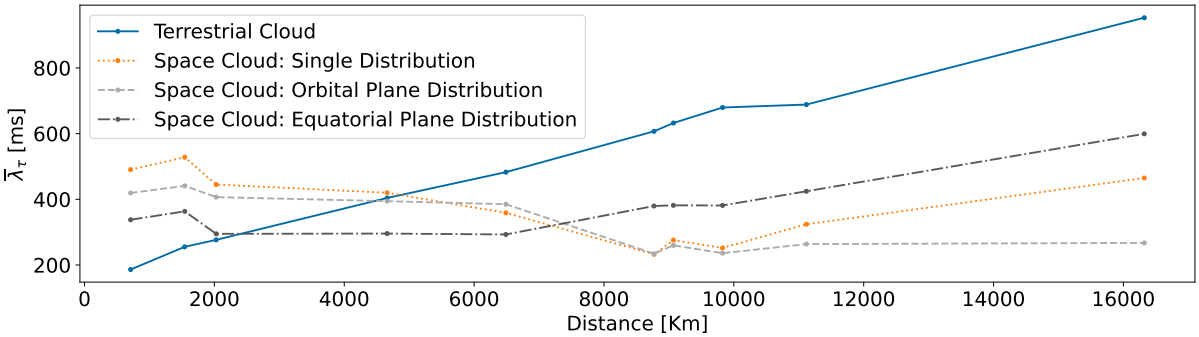

Figure A.9: Average network latency ($\bar{\lambda}_\tau$) versus geodesic distance for Terrestrial Cloud and Space Cloud configurations. In the Terrestrial Cloud, latency increases with distance due to the need to route computation requests through the satellite network to ground-based servers on the Internet. In contrast, Space Cloud configurations, with servers hosted directly on satellites, maintain stable latency regardless of distance, outperforming the Terrestrial Cloud beyond 5000 km. (From [Roj+25])

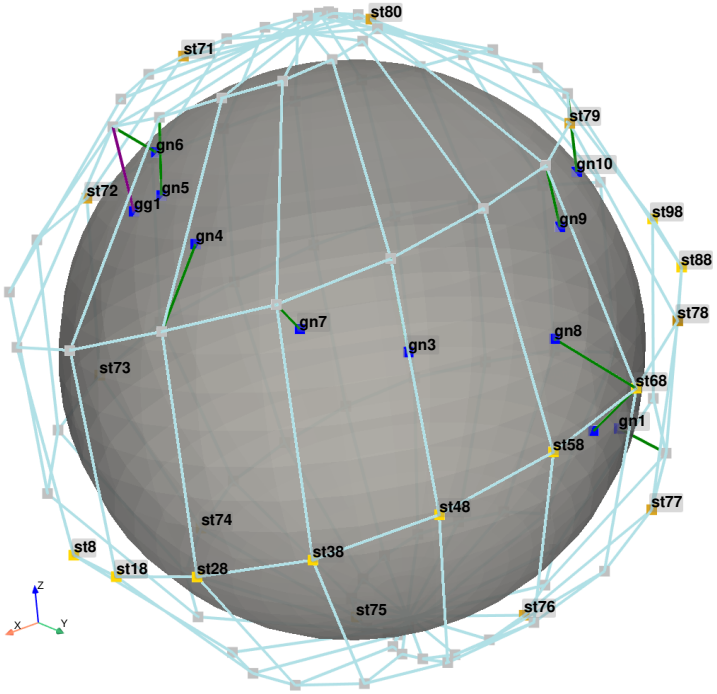

Figure A.10: Visualization of the Equatorial Plane Space Cloud configuration. Blue nodes (gn) represent ground devices requesting compute services. Yellow nodes (yellow sn) are space cloud servers deployed in the equatorial plane, while gray nodes (gray sn) act as relay satellites to route data within the constellation. (From [Roj+25])

A.2 Architectures

Architectures tackle the core research question: *How to use space?* Answering this requires thoroughly re-evaluating existing Internet architectures and assumptions to adapt them to the delays and disruptions inherent to the harsh space environment. In some cases, precise modifications to Internet protocols suffice, while more extreme challenges demand entirely novel yet backward-compatible architectures.

Contributions are divided into two key architectural domains: **Near-Earth Networks**, focusing on satellite systems operating below the geostationary belt at approximately 35,000 km altitude, and **Deep-Space Networks**, which aim to develop network architectures capable of spanning the entire Solar System and beyond.

A.2.1 Near-Earth Networks

Existing network technologies face significant limitations in remote and underserved regions, where terrestrial infrastructure is often unavailable or economically infeasible. This necessitates innovative solutions to adapt network architectures and protocols to meet the unique requirements of space-based environments.

> **Near-Earth Networks** encompass nodes located on the ground and in very low, low, medium, and geosynchronous orbits. These networks experience frequent *disruptions* and higher *delays* than terrestrial systems, yet they remain manageable with tailored adaptations of the classical Internet architecture.

My contributions to near-Earth architectures are divided into two categories. The first involves **High-Altitude Platforms**, which are stratospheric stations operating at approximately 20 km altitude and serve as a bridge between terrestrial and satellite networks. The second focuses on **Direct-to-Satellite IoT** systems, which utilize LPWAN architectures where gateways or base stations are hosted on LEO satellites, allowing devices on Earth to connect directly to the satellite.

A.2.1.1 High-Altitude Platforms

The paper in [Mad+24], developed in cooperation with Pablo Madoery and Benoit Coeugnet, a Master student I supervise, tackles the challenges of routing data in all-optical LEO constellations enhanced by HAPS, which serve as intermediaries to mitigate FSO link disruptions caused by adverse weather conditions. The resulting architecture is depicted in Figure A.11. Despite unpredictable cloud cover, the primary issue is ensuring reliable and efficient data delivery between satellites and ground stations. The study

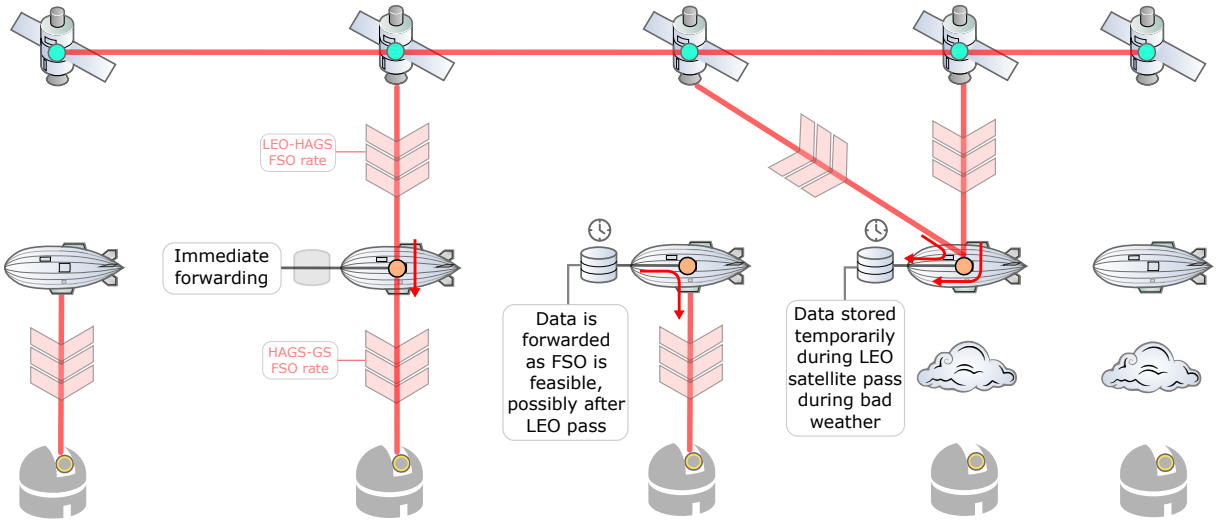

Figure A.11: The proposed High-Altitude Platform Stations (HAPS)-based architecture utilizes HAPS as temporary buffers during weather-induced Free-Space Optical (FSO) link outages, eliminating the need for backup RF links—a novel approach for FSO-based constellations. (From [Mad+24])

proposes and evaluates two novel routing solutions—*Equal Routing* and *Weighted Routing*—that leverage deterministic contact plans and integrate weather data to optimize bundle distribution across HAGS. Simulations demonstrate that Weighted Routing outperforms Greedy and Equal Routing by achieving higher delivery ratios and lower delays under varying weather conditions. The proposed solutions also enhance buffer management at HAPS. Integrating HAPS into near-Earth networks has transformative potential. To advance this frontier, I am expanding my collaborative efforts with Prof. Halim Yanikomeroglu from Carleton University. This weather-dependent architecture is well-connected with the routing under uncertain contact plans discussed in Section C.2.2. To date, two joint journal papers about this topic are under review in high-impact journals.

A.2.1.2 Direct-to-Satellite IoT

Direct-to-Satellite Internet of Things (DtS-IoT) (DtS-IoT) has been a central focus of my research career. Notably, I introduced the term DtS-IoT in [FCA19], which later became the cornerstone of the integration project during my current position at Inria.

DtS-IoT Survey This early paper [FCA19] surveyed the challenges of enabling seamless DtS-IoT connectivity, focusing on LPWAN technologies operating over LEO satellite constellations. The core finding was that IoT protocols and satellite communication standards are not designed for the unique requirements of DtS-IoT, such as managing

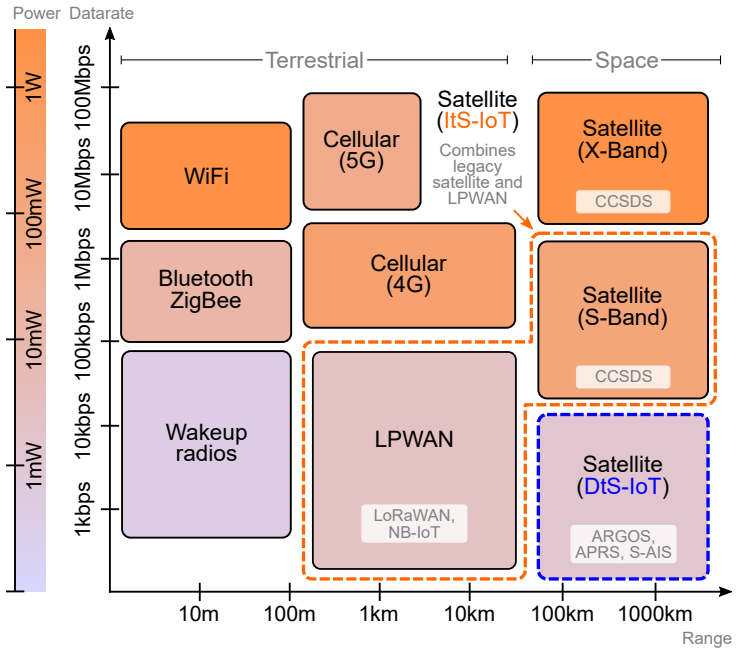

Figure A.12: DtS-IoT Protocol ecosystem across range, power consumption, and approximated data rate. This visualization highlights the positioning and trade-offs of various protocols, showcasing their suitability for satellite-based IoT systems' unique constraints and requirements. (Latest version from [FIV22])

long distances, Doppler effects, and efficient low-power communication. The paper identifies existing technology gaps and suggests future research directions, such as adapting LPWAN protocols for DtS-IoT. Figure A.12 illustrates the protocol landscape tailored for DtS-IoT, contextualized across communication range, power consumption, and data rate dimensions. This figure demonstrates that DtS-IoT occupies a distinct position in the power and range spectrum, emphasizing that no current IoT technology fully meets its specific requirements, thereby underscoring the need for tailored solutions.

DtS-IoT Architecture A subsequent landmark DtS-IoT paper in [FIV22] addressed the first definition of an architecture for the global IoT vision coined STEREO. Figure A.13 illustrates the architecture, encompassing direct and indirect satellite access. In the indirect access model, terrestrial gateways function as concentrators, aggregating data from IoT devices before relaying it to the satellite. The paper evaluates NB-IoT and LoRa/LoRaWAN as key enablers for satellite IoT, discussing their feasibility and challenges for deployment in LEO constellations. It further identifies and analyzes critical research challenges, including protocol adaptations, network function virtualization, and mobility management for IoT devices interacting with satellites. The STEREO architecture is a foundational step toward achieving DtS-IoT. STEREO gave the name to STEREO, my current ANR PRCE project, funding most of my Satellite IoT-related research at Inria.

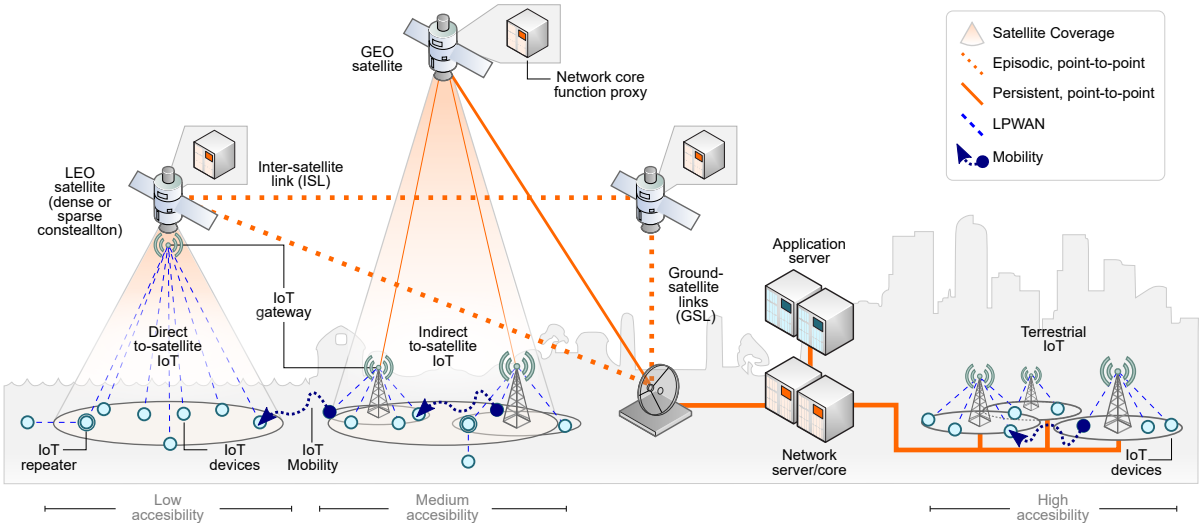

Figure A.13: The Space-Terrestrial Integrated Internet of Things (STEREO) architecture integrates terrestrial and satellite IoT networks. The system supports direct-to-satellite IoT (low accessibility regions) and indirect-to-satellite IoT (medium accessibility regions), where terrestrial gateways act as data concentrators. Persistent links through GEO satellites and episodic links via LEO satellites ensure seamless communication. (From [FIV22])

A.2.1.3 Tooling: FLoRaSat

Real-world testing for DtS-IoT is prohibitively expensive, and existing simulators fail to provide realistic orbital mechanics, inter-satellite link modeling, or end-to-end network layer integration. To this end, I introduced FLoRaSat[2].

The papers in [Fra+22b] and [Fra+22a] were the first to address the challenges of evaluating DtS-IoT systems, particularly under the constraints of long transmission distances, short satellite coverage times, and highly resource-limited IoT devices. To bridge this gap, with support from Mehdi Ait Mesbah, a Master's student I supervised, we presented FLoRaSat, an open-source, event-driven simulation tool based on OMNET++. FLoRaSat integrates realistic satellite dynamics, LoRaWAN protocols, inter-satellite links, and a modular architecture to simulate end-to-end DtS-IoT systems. The core FLoRaSat modules are depicted in Figure A.14. FLoRaSat enables studying radio access, inter-satellite routing, and network performance metrics under varying traffic loads and configurations. Case studies in [Fra+22a] identify non-intuitive phenomena such as downlink packet arrival delays (DPAD) and decoding inefficiencies in multi-gateway setups.

A recent paper [Cho+25], led by my PhD student Alexander Ylnner from the STEREO ANR project and master's student Robin Ohs from Saarland University, tackled the challenge of extending FloRaSat to support large-scale LEO satellite constellations with ISL

[2]FLoRaSat public repository: `https://gitlab.inria.fr/jfraire/florasat`.

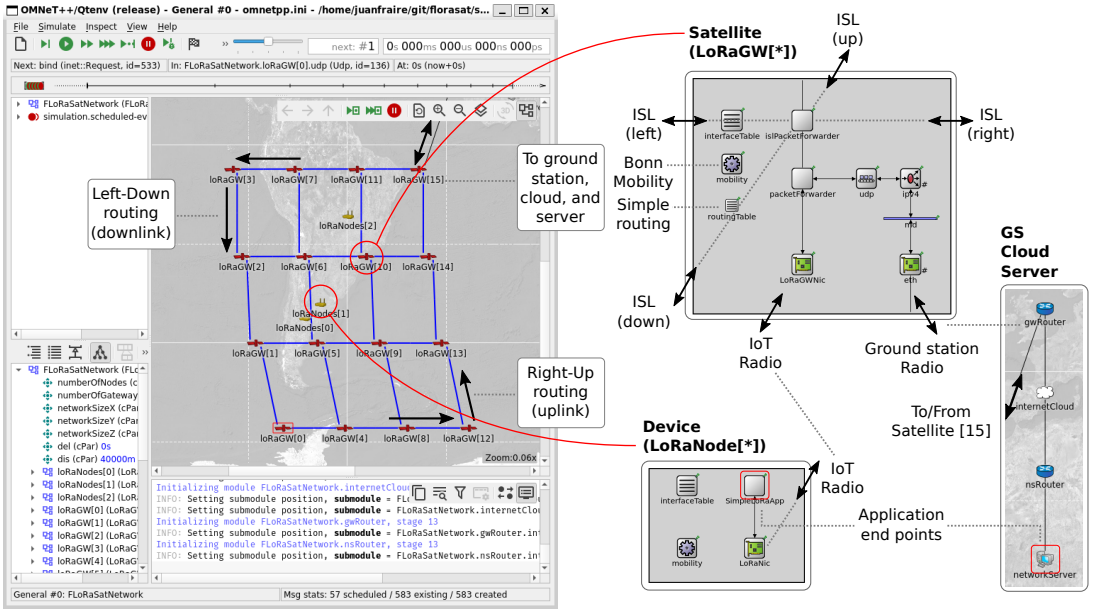

Figure A.14: FLoRaSat simulation environment and architecture. The left panel shows the OMNET++ interface, illustrating the grid-shaped LEO satellite constellation and routing paths. The right panel details the modular architecture, including LoRa-based IoT devices, satellite gateways with Inter-Satellite Links (ISL), and ground station connectivity. Data flows from IoT devices to satellites and onward to the ground station cloud server, showcasing the end-to-end DtS-IoT communication process. (From [Fra+22b])

capabilities. FloRaSat 2 incorporates advanced dynamic topology control, ISL routing sandbox capabilities for flexible routing algorithm implementation, and a visualization and analytics module. The tool enables the study of LEO constellations, such as Starlink and Iridium, and supports comparative evaluation of routing algorithms, including DDRA and DiscoRoute. Figure A.15 illustrates a sample output of FloRaSat 2. Simulations in [Cho+25] reveal the advantages of advanced routing algorithms like DiscoRoute in reducing delays and improving scalability.

Educational Impact Leveraging my extensive experience with OMNeT++ and my background in lecturing on networking topics in computer science courses, I developed concrete proposals for integrating network simulation tools into education. The paper in [FD21] tackles the challenge of modernizing undergraduate networking courses to meet the evolving demands of the Future Internet, addressing limitations in current hands-on curricula constrained by vendor-specific technologies and lacking emphasis on foundational networking concepts. The proposed methodology adopts an action-research approach, embedding active and creative learning techniques within a redesigned curriculum for the networking course at FAMAF, UNC. The methodology comprises two stages: development, which focuses on algorithm design and coding using Python-based networking applications, and analysis, which centers on evaluating

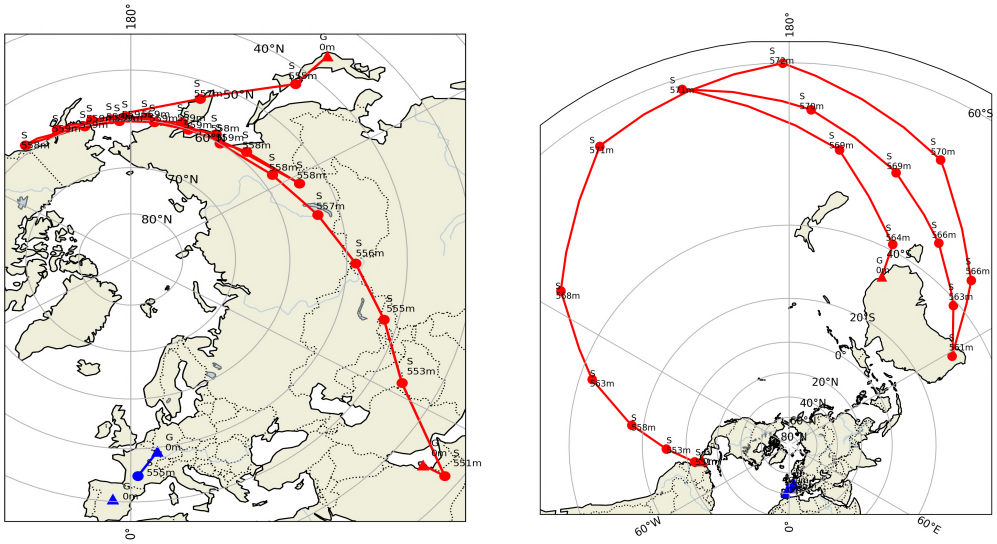

Figure A.15: Comparison of routing paths in FloRaSat 2 for a message transmitted across a LEO satellite constellation using a) DDRA and b) DiscoRoute algorithms. DDRA follows a shortest-path approach, prioritizing minimal hops but resulting in increased propagation delays. DiscoRoute leverages an optimized path selection, minimizing overall latency by balancing link delays. (From [Cho+25])

routing and flow control algorithms via the OMNeT++ simulation platform. Innovative teaching strategies include game-based learning, optional starred tasks for advanced exploration, and public expositions to drive student engagement and computational thinking. The results reveal substantial improvements, with 92% of students achieving course objectives in the revised cycle—a 7% increase compared to the baseline—while significantly enhancing motivation and technical proficiency in core networking concepts.

A.2.2 Deep-Space Networks

While my PhD research focused on CONAE's SARE mission, a near-Earth network, its sparse topology and episodic ISL communications closely resembled the dynamics of deep-space networks. These networks share key challenges such as planetary occlusion, resource constraints, and limited architectures, distinguishing them from traditional satellite constellations like Starlink. Consequently, my contributions were naturally adaptable to interplanetary and solar system-scale networks, extending their relevance and applicability to deeper space contexts.

Deep-Space Networks extend beyond geosynchronous orbits, encompassing nodes on the Moon, other planets, and beyond. These networks face prolonged *disruptions* caused by planetary occlusion and extreme *delays*, ranging from minutes to hours, due to the finite speed of light over vast distances. Such challenges necessitate fundamentally different network architectures.

This section highlights my contributions to Deep-Space Network architectures, which extend beyond geosynchronous orbits and encompass nodes on the Moon, other planets, and interplanetary contexts. Central to these efforts is the development and application of DTN principles, as explored in my comprehensive DTNBook [FFB17a], a foundation for much of the work presented here. This section showcases key contributions in designing scalable, robust networking systems for the Solar System and beyond, from **Interplanetary** architectures to **Interstellar** connectivity, derived **Toolings**, and actual **Protocol Stack** implementations.

A.2.2.1 Interplanetary

I am an active contributor in the interplanetary networking domain. I serve as a board member of the Interplanetary Networking Special Interest Group (IPNSIG), promoting the adoption of interplanetary protocols across academia and industry globally. Additionally, I am an academic stakeholder in the Dynamic Coalition on the Interplanetary Internet under the Internet Governance Forum (IGF). Furthermore, at the time of writing, I am actively involved in ESA's Pre-Phase A Solar System Internet (SSI) project, which aims to lay the groundwork for an interplanetary-scale network, positioning Europe as a key stakeholder in this strategic domain. Figure A.16 presents the key aspects and actors of the project and early analysis results about the propagation latency distribution across the planets of the Solar System.

While much of the research discussed in the remainder of this section—particularly on design and operations in Section B, and routing and transport in Section C—applies to interplanetary networking, this section highlights architecture-specific contributions.

Mars DtS-IoT Architecture In collaboration with Oscar Ledesma and Paula Lamo, from UNIR, we contributed with [Led+24]. This paper presents an architectural framework for DtS-IoT communication on Mars, considering the planet's unique environmental constraints and the need for efficient data transmission between surface sensors, orbiting satellites, and Earth. We propose an architectural framework that adapts DTN and Long-Range (LoRa) technologies. The architecture comprises multiple layers, including Martian IoT nodes, satellite linkage, constellation configuration, and Earth com-

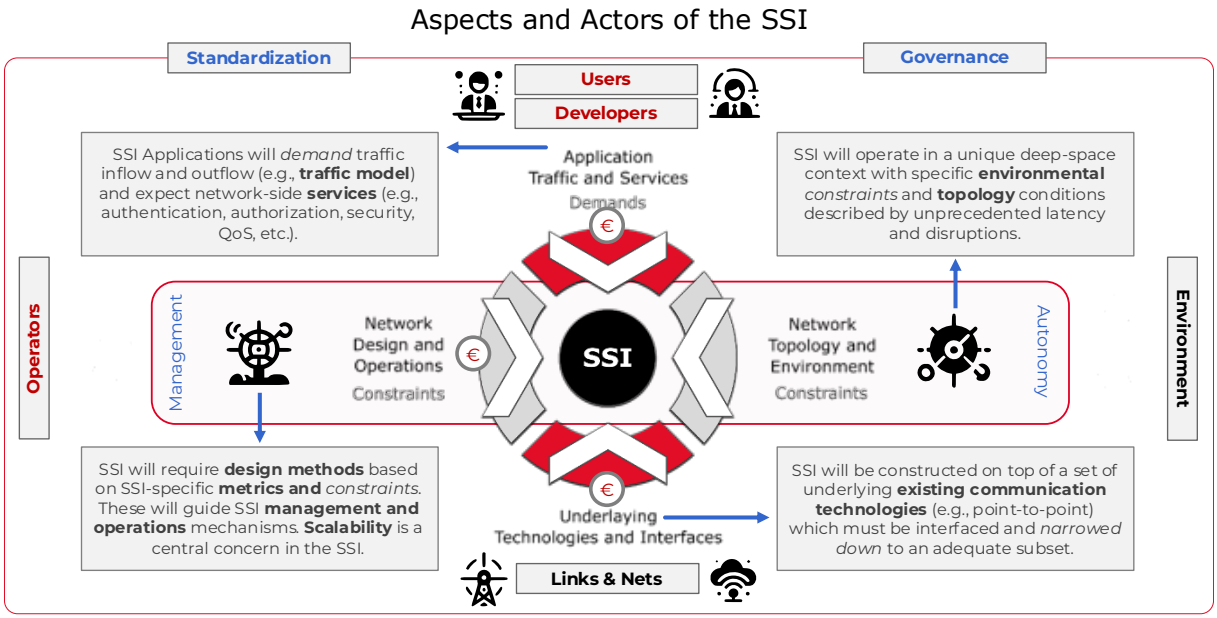

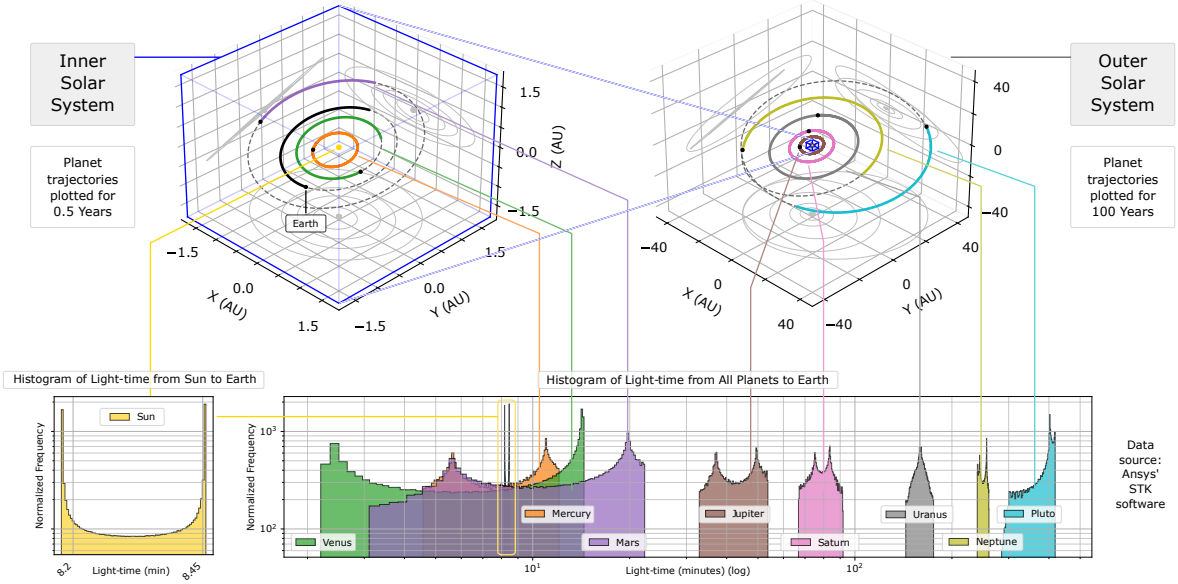

Figure A.16: ESA's Pre-Phase A Solar System Internet (SSI) project. On the top, the diagram highlights the critical elements of the SSI, including user and developer demands for traffic and services, environmental constraints unique to deep-space contexts, and the need for scalable network design and operations. It emphasizes standardization, management, governance, and autonomy, illustrating how existing communication technologies (e.g., point-to-point links) must be integrated and tailored for the unprecedented challenges of interplanetary networking. On the bottom, a visualization of planetary trajectories and light-time latency histograms within the inner and outer Solar System. The top panels display planet trajectories plotted for 0.5 years (inner Solar System) and 100 years (outer Solar System) in a 3D coordinate system. The bottom panels present histograms of light-time from the Sun to Earth and from each planet to Earth. Data was generated using Ansys' STK software. (Courtesy of ESA)

munication. The feasibility analysis focuses on four critical aspects: maximum communication range, satellite constellation's message processing capacity, communication frequency and visibility based on constellation arrangement, and interplanetary data transmission capabilities of LoRa-based IoT devices. The scenario and key visibility metrics derived from this analysis are shown in Figure A.17. We are extending FloRaSat (see Section A.2.1.3) to enable finer-grain simulation and analysis of this Martian IoT setup.

A.2.2.2 Tooling: IPN-V

This paper in [Le +24] addresses the challenge of visualizing and planning interplanetary communication networks, characterized by dynamic topologies, time-varying latencies, and disruptions due to planetary occlusions. The proposed solution, the Interplanetary Network Visualizer (IPN-V)[3], bridges astrodynamics and network engineering by enabling realistic modeling and visualization of interplanetary DTN scenarios. The tool, built on Unity 3D and powered by modular components like an IPN-D generation module, supports input from Keplerian models. It provides interactive 3D visualizations of network dynamics, including contact plans compatible with NASA's ION and HDTN formats. Key outcomes include a versatile platform for education, outreach, and operational analysis, validated through its capability to visualize up to thousands of simultaneous contacts with minimal computational overhead. Notably, IPN-V was used to render the IPNSIG 100-Year Vision Video for the SSI, that won a contest at the IPNSIG. Figure A.18 presents a series of screenshots of the visualization capabilities of IPN-V. The IPN-V's user-friendly interface and robust architecture position it as a critical tool for advancing interplanetary networking research and collaboration.

The work on IPN-V was led by Alice Le-Bihan, a Master's student I supervised at INSA Lyon during her first internship. Alice subsequently completed a second internship at ESA, under the supervision of Felix Flentge, who also leads the aforementioned SSI project in which I am actively involved. During her ESA internship, Alice contributed to developing custody transfer mechanisms for DTN, where the intuitive insights provided by IPN-V played a pivotal role. At the time of writing, Alice is preparing for a co-funded PhD thesis between ESA and Inria, focusing on developing a Digital Twin for the SSI. Felix Flentge and I will co-supervise her PhD, with informal involvement from Bob Short from NASA Glenn Research Center, ensuring the Digital Twin is applicable and beneficial for both NASA and ESA. The title of the proposal is *Advancing the Solar System Internet Through Digital Twin-Driven Delay-Tolerant Network Optimization and Operations*, and the details can be found in this PhD project document.

[3]IPN-V public repository: `https://gitlab.inria.fr/jfraire/ipn-v`

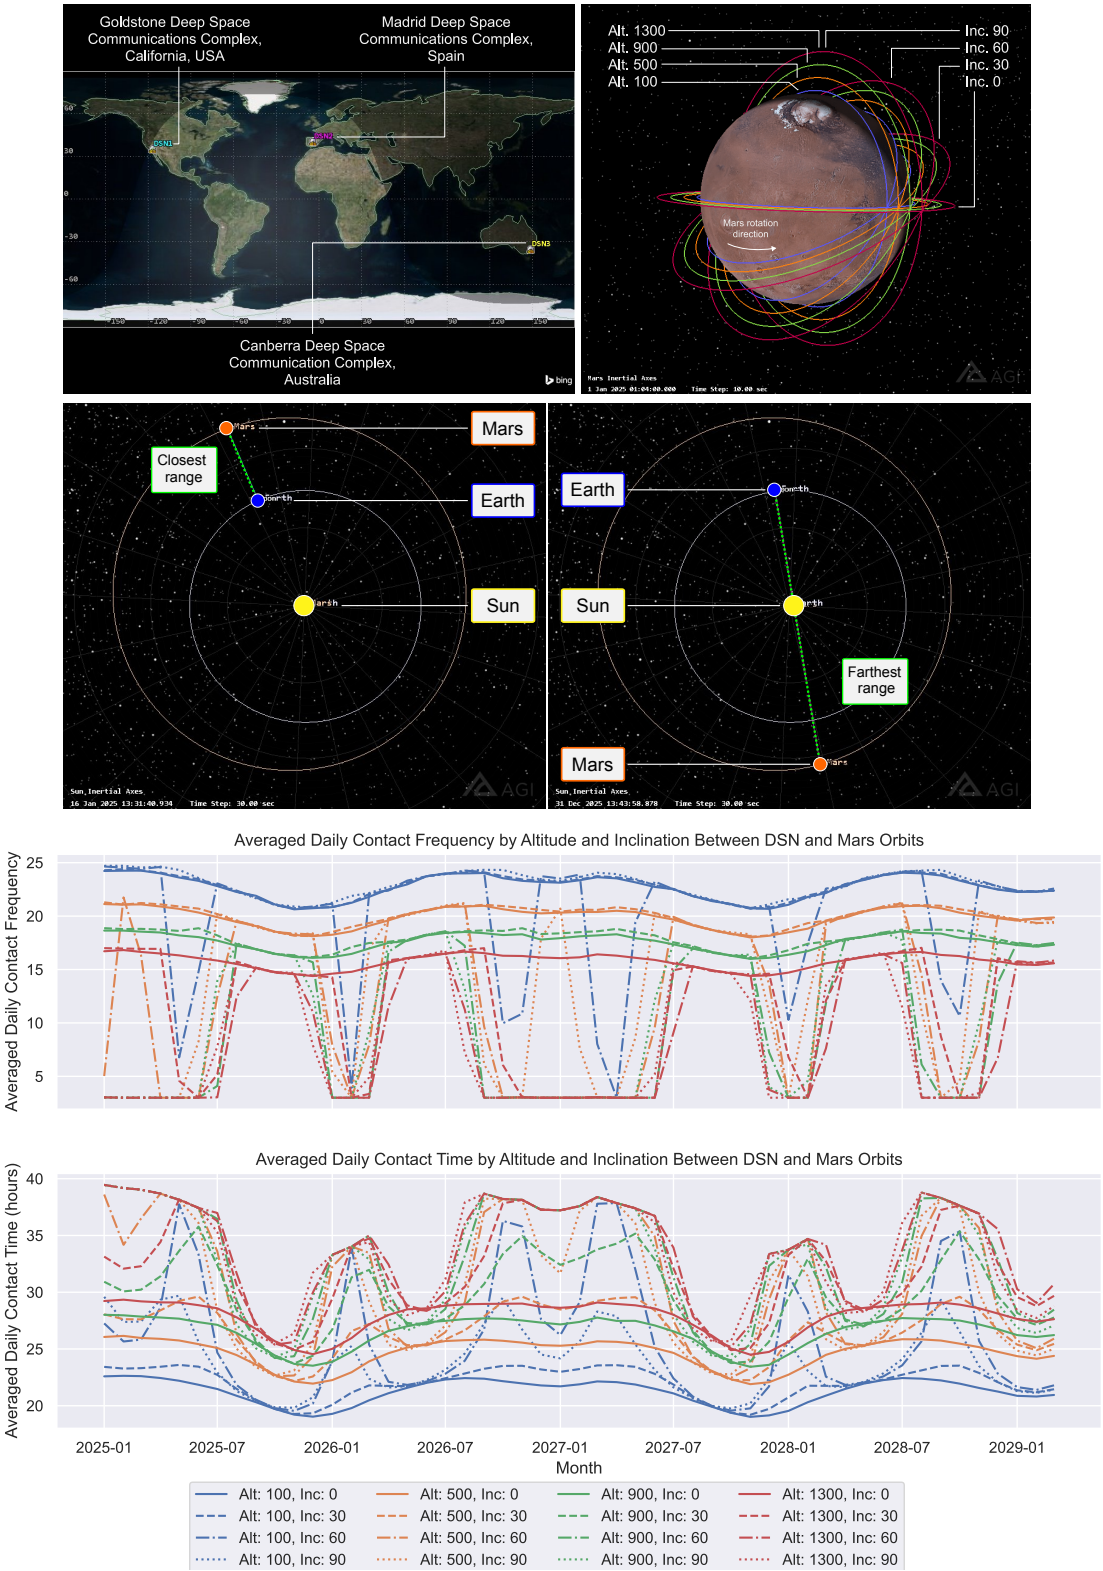

Figure A.17: Mars DtS-IoT architecture analysis. The top section illustrates the ground stations on Earth, along with the altitudes and inclinations of Mars orbiters. The center section highlights Earth and Mars's closest and farthest alignment scenarios. The bottom section presents the daily contact duration and frequency between Earth stations and Mars orbiters. (From [Led+24])

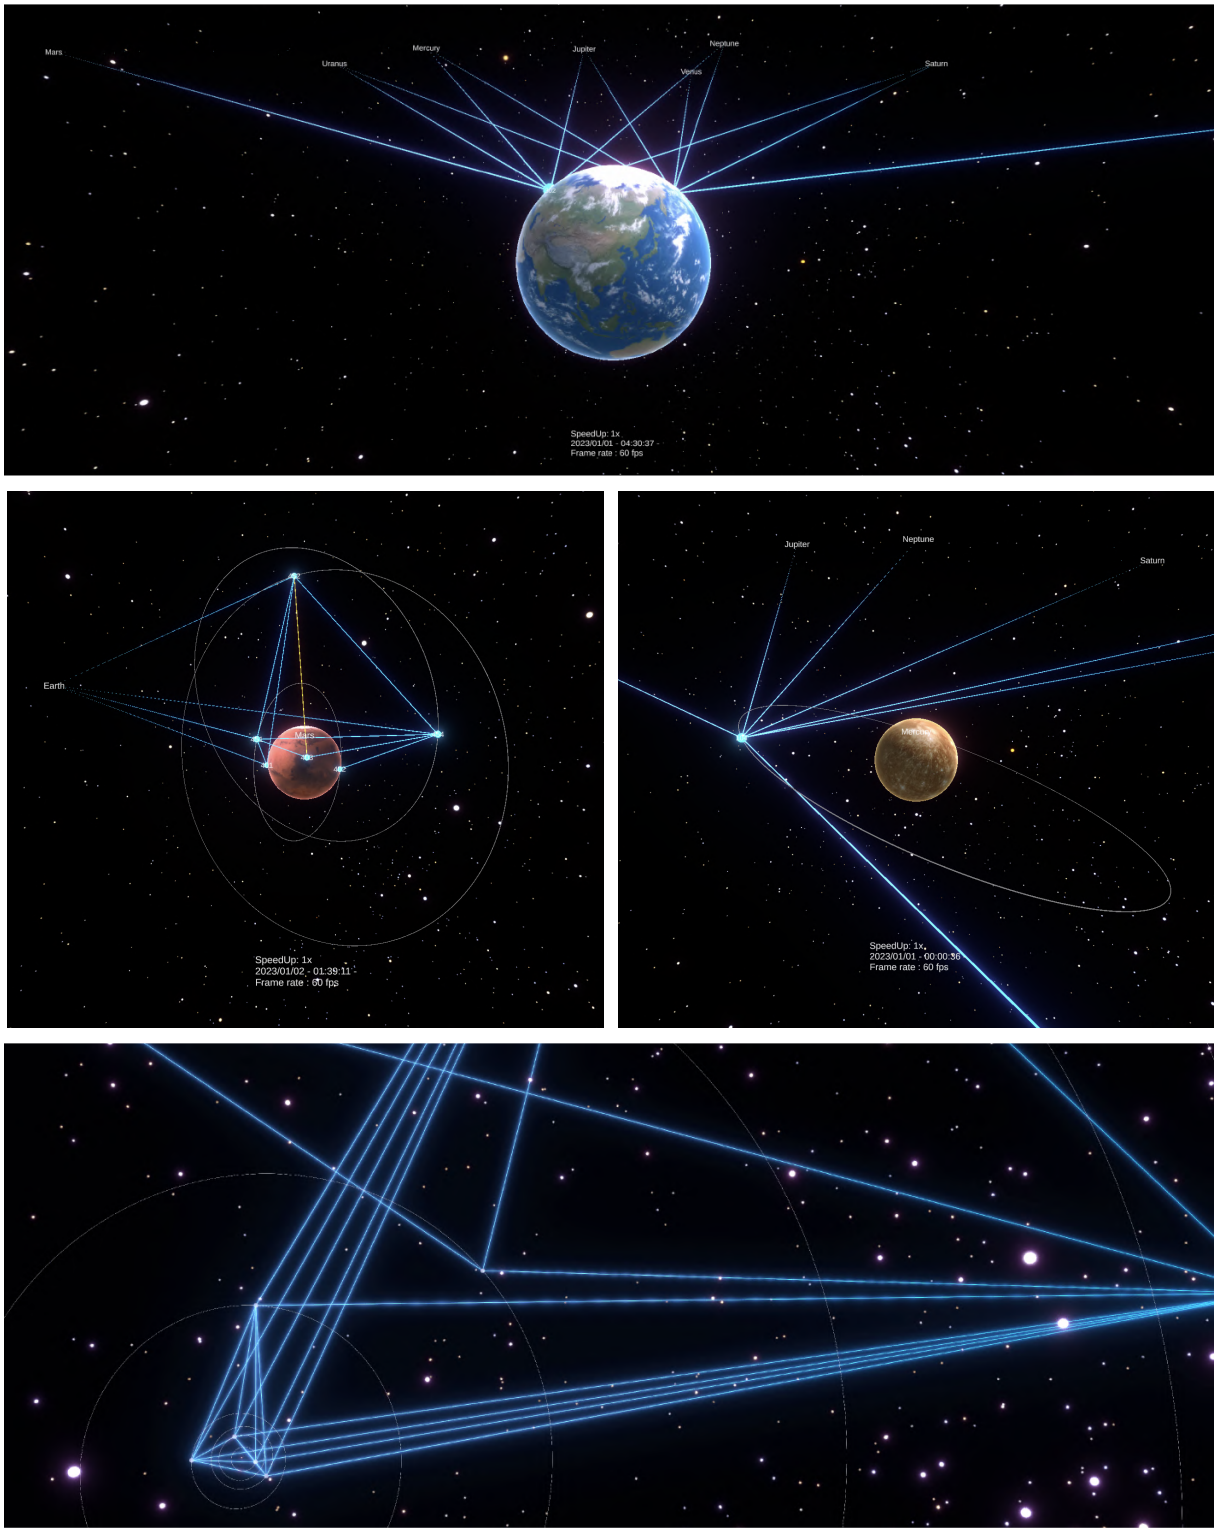

Figure A.18: Screenshots of the Interplanetary Network Visualizer (IPN-V) showcasing dynamic visualizations of interplanetary communication networks. The top panel highlights Earth-centric communication paths to multiple celestial bodies, including Mars, Jupiter, and Saturn. The middle panels focus on Mars, illustrating bidirectional links with Earth and neighboring planets. The bottom panel provides a detailed view of the intricate link topology within a segment of the interplanetary network. (From [Le +24])

A.2.2.3 Stacks: ION and μD3TN

While my core research expertise lies in developing methods and algorithms for space networking, I have successfully translated this knowledge into implementing operational-grade deep-space protocol stacks throughout my career. Notably, I have contributed to the Interplanetary Overlay Network (ION), developed by NASA's Jet Propulsion Laboratory (JPL), and μD3TN, a state-of-the-art DTN protocol stack developed and maintained by D3TN GmbH, a company based in Dresden, Germany. My collaboration with D3TN is further strengthened through our joint participation in the MISSION.

ION An early paper during my PhD in [Alf+14] addresses the challenge of extending DTN protocols to intra-spacecraft communication by enabling BP data transmission over SpaceWire, a standard for onboard spacecraft networking. In this work, I supervised Marcos Alfonzo, a space software developer from CONAE, to develop a Convergence Layer Adapter (CLA) for SpaceWire within NASA's ION software[4]. ION is a software implementation of DTN protocol suite. It enables reliable data transfer across terrestrial, interplanetary, and deep-space networks by employing store-and-forward mechanisms, custody transfers, and BP functionalities. Our CLA allows the BP to interface with SpaceWire, ensuring compatibility with the SpaceWire protocol's high-speed, low-power, and full-duplex characteristics. The implementation is evaluated using a LEON3FT-RTAX-based onboard computer test bench, comparing its performance against Ethernet-based CLAs in the ION stack. Results demonstrate that the SpaceWire CLA significantly outperforms its Ethernet counterparts, achieving an order of magnitude higher throughput for DTN payloads with minimal additional overhead. Figure A.19 illustrates the architecture of the SpaceWire CLA implementation on the LEON3FT-RTAX FPGA alongside the results from the evaluation campaign. The figure highlights the integration of the SpaceWire CLA within the ION protocol stack and showcases its performance metrics compared to alternative CLAs such as UDP and TCP. This work establishes a foundational step for integrating DTN into intra-spacecraft domains, paving the way for advanced spacecraft network architectures.

μD3TN This paper in [Wal+24] addresses the challenge of developing a flexible, efficient, and modular DTN protocol stack tailored for diverse use cases, from terrestrial to interplanetary networks. To tackle this, I collaborated with D3TN GmbH in the context of the MISSION to present the architectural evolution of μD3TN, transitioning from its origins as μPCN for microcontrollers to a sophisticated, software-defined DTN protocol stack[5]. Key innovations include platform independence, concurrent support for

[4] ION public repository: https://github.com/nasa-jpl/ION-DTN
[5] μD3TN public repository: https://gitlab.com/d3tn/ud3tn

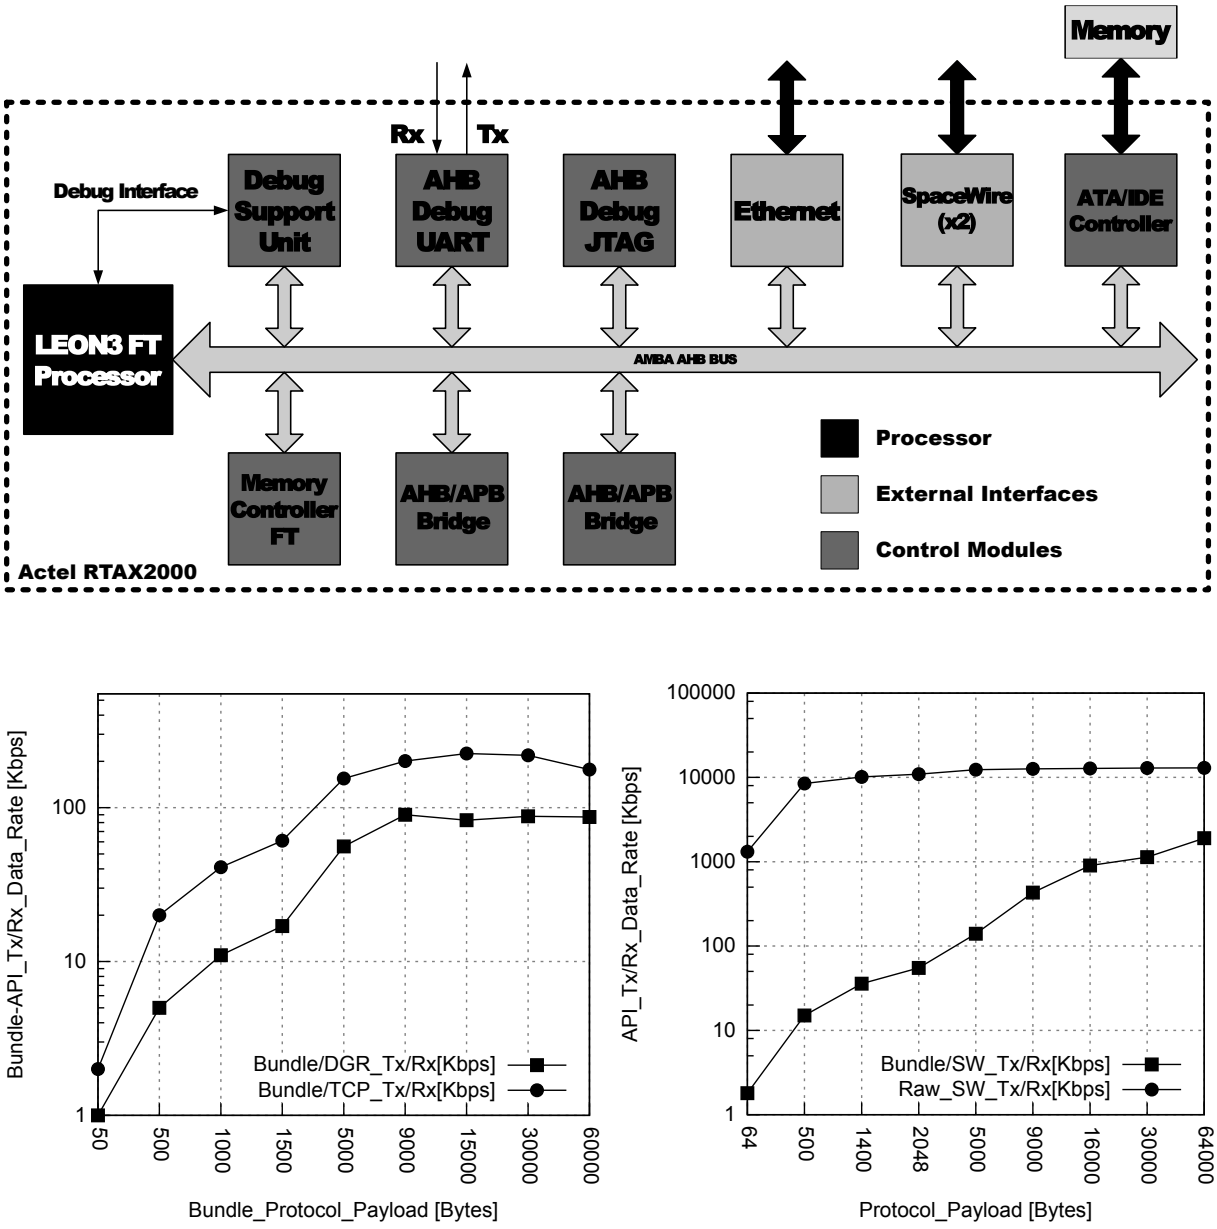

Figure A.19: Architecture and performance evaluation of the SpaceWire Convergence Layer Adapter (CLA) implemented on the LEON3FT-RTAX FPGA. The top panel illustrates the internal architecture, showing the integration of the SpaceWire CLA with the LEON3FT processor and associated interfaces, including Ethernet and memory controllers, via the AMBA AHB bus. The bottom panels present performance metrics: data transmission and reception rates (Kbps) and drop rates (%) for various Bundle Protocol (BP) payload sizes. These results compare Ethernet-based implementations (DGR and TCP CLAs already present in ION distribution on the left) and SpaceWire (SW, our own CLA implementation on the right), highlighting the superior throughput and reduced overhead achieved by the SpaceWire CLA for DTN applications. This development was delivered to CONAE as part of the SARE distributed satellite mission in LEO. (From [Alf+14])

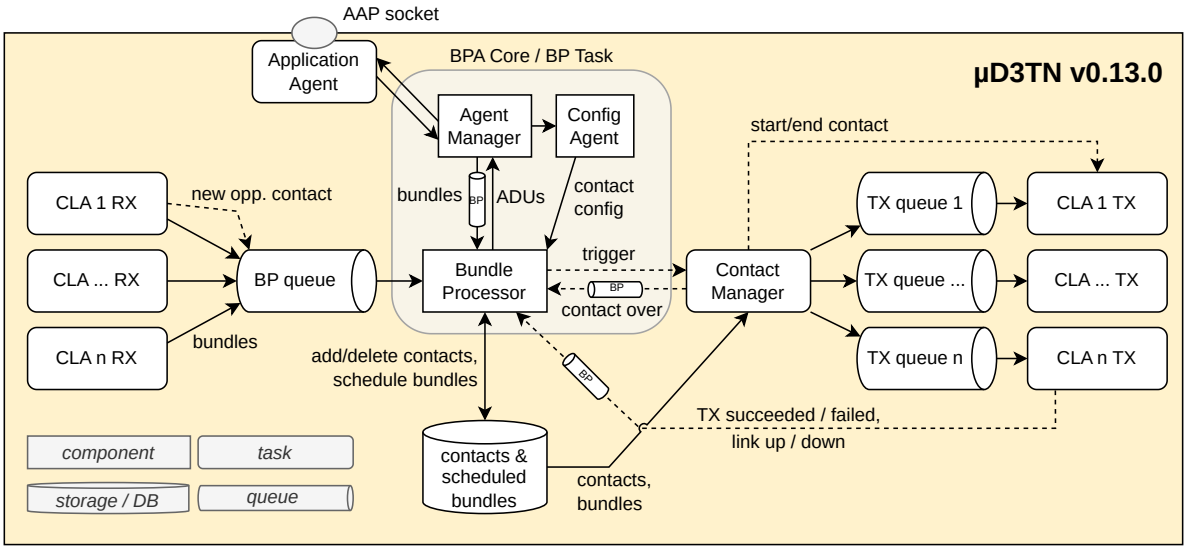

Figure A.20: Architecture of μD3TN v0.13.0. This modular DTN protocol stack comprises several key components and processes. The Bundle Processor (BP) serves as the core of the Bundle Protocol Agent (BPA), handling incoming and outgoing bundles, managing Application Data Units (ADUs), and interfacing with the Agent Manager and Config Agent. The Application Agent, connected via an AAP socket, enables communication between the BPA and external applications, facilitating efficient data exchange through the Application Agent Protocol (AAP). The Contact Manager is critical in scheduling and tracking contacts, triggering transmission events, and managing link status such as start/end contacts or link up/down while ensuring proper interaction with TX queues and handling successful transmissions or failures. Convergence Layer Adapters (CLAs) support data flows between the Bundle Processor and the physical interfaces, with distinct RX (receive) and TX (transmit) queues for each layer, enabling compatibility with diverse communication channels. Persistent storage maintains contact configurations and scheduled bundles, ensuring reliability and scalability across intermittent network conditions. (From [Wal+24])

BPv6 and BPv7, a generic convergence-layer adapter (CLA) interface, and the Application Agent Protocol (AAP) for enhanced application-layer interaction. Further enhancements, such as Bundle-in-Bundle Encapsulation (BIBE), the Generic Bundle Forwarding Interface, and persistent storage mechanisms, are integrated to improve scalability, routing flexibility, and cross-platform compatibility. Results from practical applications, including ESA's OPS-SAT mission, validate μD3TN's robustness and highlight its potential for scalable DTN deployments across terrestrial and space domains. This refined architecture, illustrated and described in Figure A.20, positions μD3TN as a state-of-the-art solution for modern DTN challenges.

A.2.2.4 Interstellar

In collaboration with Prof. Elena Fantino, the paper in [Fra+19] addresses the ambitious challenge of enabling reliable interstellar communication with the TRAPPIST-1 system, located 39.5 light-years away, where extreme delays and disruptions render traditional Internet protocols ineffective. I was the first to propose leveraging DTN principles to build a multi-hop relay network that employs the store-and-forward paradigm and custody transfer mechanisms to ensure reliable data delivery. We evaluated the trade-offs between relay spacing, buffer sizes, and achievable data rates using link budget analysis and retransmission delay modeling. Key outcomes include demonstrating the feasibility of deploying thousands of relays spaced approximately one light day apart to achieve manageable end-to-end delays and buffer requirements. Figure A.21 provides a 3D visualization of the interstellar domain, showcasing nearby stars and an interstellar DTN path connecting our Solar System to the TRAPPIST-1 system. The study underscores the critical role of DTN protocols in enabling interstellar networking while highlighting the technological challenges and generational efforts required to build such an ambitious infrastructure.

A.3 Environment

This section focuses on understanding the space environment, with a particular emphasis on the radiation effects on electronic devices critical to the operation of space networks. While other environmental challenges, such as extreme temperature fluctuations, outgassing in a vacuum, and mechanical vibrations during launch, are significant, they are addressed by dedicated disciplines like thermal control, structural engineering, and harness design, which are less directly connected to communications and networking. This is not the case with the effects of charged particles.

Radiation effects on microelectronics directly challenge the integrity of networking and communication systems. These effects, including Single-Event Effects (SEE) and Total Ionizing Dose (TID) breakdowns, can severely impact the performance of modems, transponders, and other link and network electronics, making this a critical area of investigation within the context of space networking.

After earning my PhD, I pursued a postdoctoral position at TIMA Laboratories, part of Université Grenoble Alpes, under the guidance of Dr. Raoul Velazco, a renowned expert in the effects of radiation on electronic systems. While my primary focus during this postdoc was on networking reliability—a topic extensively discussed in Sections C.2.1.2 and C.2.2—I also had the opportunity to collaborate with the team on

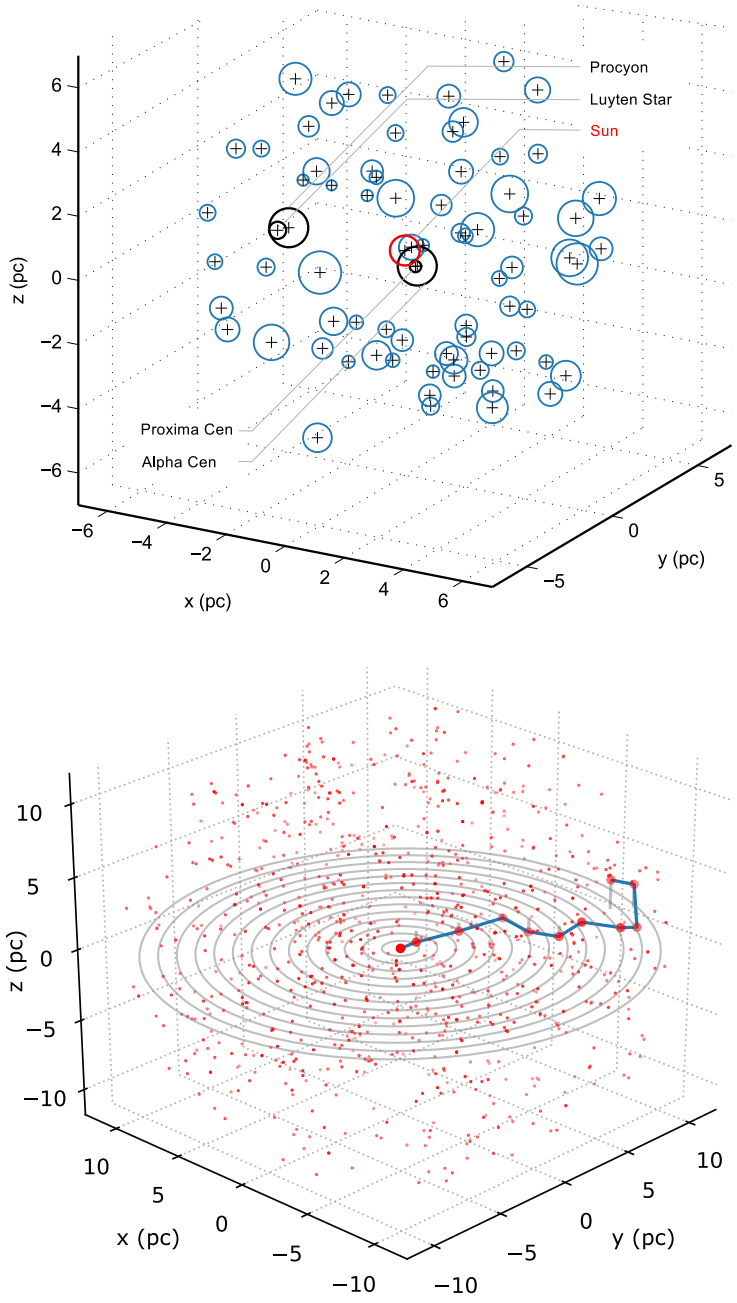

Figure A.21: Interstellar networking. On the top, a 3D representation of the Solar Neighborhood, showing nearby stellar systems within 6 parsecs (pc) of the Sun. Key stars such as Proxima Centauri, Alpha Centauri, Luyten's Star, and Procyon are labeled. Each marker's size corresponds to the stars' region of influence, and the Sun is highlighted in red. On the bottom, another 3D map extending to 10 parsecs from the Sun illustrates a broader view of nearby stars and their distribution within the galaxy. The multi-hop overlay represents potential communication pathways to the TRAPPIST-1 system, a main attraction for interstellar networking among the community due to the presence of habitable planets.(From [Fra+19])

radiation experiments and device characterization under harsh space environments. The insights gained from these efforts deepened our understanding of failure models and informed the development of accurate abstractions, which underpin the reliability strategies detailed in the sections above.

This section summarizes my work's key contributions and outcomes in this area. After a short overview of the radiation effects, I summarize my works in **radiation exposure campaigns**, **fault injection campaigns**, and **networks on chip applications**.

Radiation Effects Overview The space radiation environment introduces critical challenges for electronic systems, particularly in satellite networks, due to Single-Event Effects (SEE) and Total Ionizing Dose (TID). SEE, caused by high-energy particle strikes, includes phenomena such as Single-Event Upsets (SEU), which are bit flips in memory or registers that can disrupt data integrity, and Single-Event Transient (SET), which are temporary voltage spikes in logic circuits that may propagate as errors if not mitigated. These effects result in transient disruptions such as bit flips or system crashes, while TID gradually degrades electronic components, reducing performance and lifespan. These radiation-induced effects can lead to permanent or transient failures at the node and link levels, severely impacting communications reliability through data loss and link disruptions.

Conventional space networking approaches are poorly equipped to address these challenges due to their reliance on deterministic contact plans, inability to adapt to real-time dynamic failures, and lack of cross-layer integration to account for physical-layer radiation effects. This gap presents a unique opportunity to develop uncertainty-aware data handling in space networking protocols. Real-time fault detection and adaptive responses could be incorporated to enhance robustness and efficiency in space communication networks. Integrating such mechanisms would fundamentally improve the reliability of satellite networks operating in harsh radiation environments.

A.3.1 Radiation Exposure Campaigns

Radiation Exposure Campaigns, often called beam exposure campaigns, involve testing electronic components or systems under controlled radiation environments to evaluate their resilience to space-like conditions. These campaigns typically use facilities like particle accelerators or heavy-ion beams to simulate the effects of cosmic radiation. Researchers can assess failure mechanisms, characterize vulnerabilities, and validate mitigation strategies by exposing devices to high-energy particles. This ensures that components can operate reliably in space's harsh radiation environment. These tests are critical for qualifying space-bound hardware and improving its robustness for

mission-critical applications. The following section summarizes my experience in this domain with a related **methodology** and a **memory characterization** campaign.

A.3.1.1 Methodology

Wireless Access for Radiation Beam Experiments Our paper in [Sán+19] presents a novel embedded wireless system designed to measure the effects of ionizing radiation on electronic devices, addressing challenges in safety, flexibility, and accessibility during radiation campaigns. The method combines the NETFI-2 fault injection tool, described in the section below, with a custom interface board and mobile wireless control and monitoring application. The system employs Commercial-Off-The-Shelf (COTS) components, including Bluetooth-enabled microcontrollers, to replace traditional wired setups. This allows experiments to be conducted remotely, improving operator safety and enabling real-time data acquisition. Validation experiments at the TANDAR facility in Buenos Aires demonstrate the system's ability to detect SEEs, such as bit flips in SRAM devices exposed to radiation. Results highlight the system's reliability and scalability in harsh environments. This innovative approach could be applied to satellite component testing and fault injection campaigns for CubeSat missions. It bridges the gap between traditional fault injection methods and modern wireless monitoring solutions, advancing radiation effect analysis for space electronics.

A.3.1.2 Memory Characterization

Our study in [Cle+17] and its journal extension in [Cle+18], both developed in cooperation with Juan Clemente, from Universidad Complutense de Madrid, we investigate the sensitivity of three generations of commercially available off-the-shelf (COTS) Static Random-Access Memory (SRAM) chips—manufactured at 130 nm, 90 nm, and 65 nm CMOS nodes—to SEUs under 14.2 MeV neutron irradiation at ultra-low bias voltages. The problem evaluates how scaling down bias voltage and technological integration impacts SEU sensitivity, especially under harsh radiation environments. The method employed involved static irradiation tests on these SRAMs at voltages ranging from 0.5 V to 3.15 V. Errors were characterized by cross-section analysis using the MUSCA-SEP3 predictive modeling tool. Key outcomes reveal that 130 nm and 65 nm SRAMs' sensitivity to SEUs increases significantly at near-threshold voltage levels. Meanwhile, 90 nm SRAMs displayed erratic behavior attributed to micro single-event latchups (micro-SELs). The study also demonstrated the effectiveness of error-correcting code (ECC) mechanisms in mitigating errors in the 65 nm SRAM, achieving over 99% correction efficiency even under high neutron fluences. Moreover, trends in SEU multiplicity and large-scale events, such as multi-bit upsets (MBUs) and single-event functional inter-

rupts (SEFIs), were detailed, with the results validated against MUSCA-SEP3 predictions. These findings underline the nuanced interplay between bias voltage, technology scaling, and SEU resilience in SRAMs, providing critical insights for their application in radiation-prone environments like avionics and aerospace.

A.3.2 Fault Injection Campaigns

Fault injection is a technique for testing and evaluating electronic systems' and components' robustness, reliability, and fault tolerance under controlled fault scenarios. By deliberately introducing errors or disturbances into a system, fault injection simulates the effects of real-world faults caused by factors like radiation, aging, or environmental interference. This helps developers identify vulnerabilities, assess failure modes, and design mitigation strategies at lower costs than radiation campaigns. The following section summarizes my experience in this domain, which includes the development of a fault injection **methodology** and the **characterization of a processing** device.

A.3.2.1 Methodology

NETFI-2 Our paper in [Sol+17] presents an advanced methodology for fault injection in FPGA-based systems, addressing key challenges in testing the susceptibility of electronic circuits to radiation-induced errors, specifically SEUs and SETs. The study tackles the limitations of existing fault injection methods, which either involve high hardware overhead, lack real-time capabilities, or depend on external platforms for fault injection control. We proposed NETFI-2, a revised methodology that integrates fault injection management directly into the FPGA using a MicroBlaze soft-core processor, eliminating the need for external controllers and minimizing resource usage. NETFI-2 employs a streamlined workflow, illustrated in Figure A.22. The Device Under Test (DUT) is synthesized into a modified netlist using the MODNET tool. The controller, implemented within the FPGA, communicates via the AXI-Lite interface to inject faults and collect results. The system allows simultaneous injection of multiple faults and supports SEU and SET testing, providing flexibility through embedded software configuration. The method was validated using a Stochastic Bayesian Machines (SBM) module (discussed below) as a DUT, demonstrating a strong correlation between simulation-based and NETFI-2-based results, with susceptibility rates of approximately 13% for SEUs and 31% for SETs. Notably, NETFI-2 achieved a 300% improvement in execution speed compared to simulation-based campaigns, completing exhaustive fault injections in under one second. The methodology significantly reduces hardware overhead, with DUT modifications requiring only 25% and 9.25% additional LUT resources for SEU and SET injection, respectively, while maintaining accuracy and scalability for larger designs. These

advancements establish NETFI-2 as a powerful tool for radiation testing and reliability evaluation of FPGA-based systems.

It is worth noting that Dardo Ariel Viñas, as part of his Master's thesis at UNC, which I co-supervised, completed the development of NETFI-1, which is now publicly accessible in https://github.com/LCSR-lab/NetFi3.

A.3.2.2 Processors Characterization

Characterization of Bayesian Machines Our paper in [Coe+16] and its journal extension in [Coe+17] addresses the challenge of assessing the robustness of hardware-based Stochastic Bayesian Machines (SBM)s in radiation-prone environments, focusing on their resistance to soft errors caused by SEUs and SETs. The proposed method implements a fault-injection campaign at the Register Transfer Level (RTL) using a non-intrusive modification of the SBM's logic design, specifically targeting flip-flops and lookup tables. We conducted two fault-injection campaigns: the first analyzes individual SBM slices, and the second evaluates the complete SBM system operating under the stochastic computing paradigm. The results demonstrate that while individual SBM slices are prone to significant error rates when analyzed deterministically, the stochastic encoding scheme mitigates the impact of soft errors in the complete SBM system. Under extreme fault-injection conditions (e.g., up to 30% fault injection probability), the deviation in output probability distributions remains below 1.3%, illustrating stochastic architectures' intrinsic fault tolerance. These findings validate SBMs as energy-efficient, fault-resilient candidates for probabilistic applications in harsh radiation environments. SBMs offer a compelling alternative to conventional deterministic hardware hardened by redundancy techniques.

A.3.3 Networks on Chip Applications

This section describes two radiation-related works in the context of Networks on Chip (NoC) applications.

Wireless DTN NoCs Our papers in [Fer+19] and [Fer+21] introduce the concept of Delay-Tolerant Networks on Chips (DTNOCs) to enhance the dependability of wireless NoC. It addresses challenges in traditional NoC, including high fault rates caused by transient or permanent errors (e.g., ionizing radiation) and the limitations of existing bus-based interconnection methods, which suffer from poor scalability, energy inefficiencies, and high propagation delays. We adapted delay-tolerant networking principles, ini-

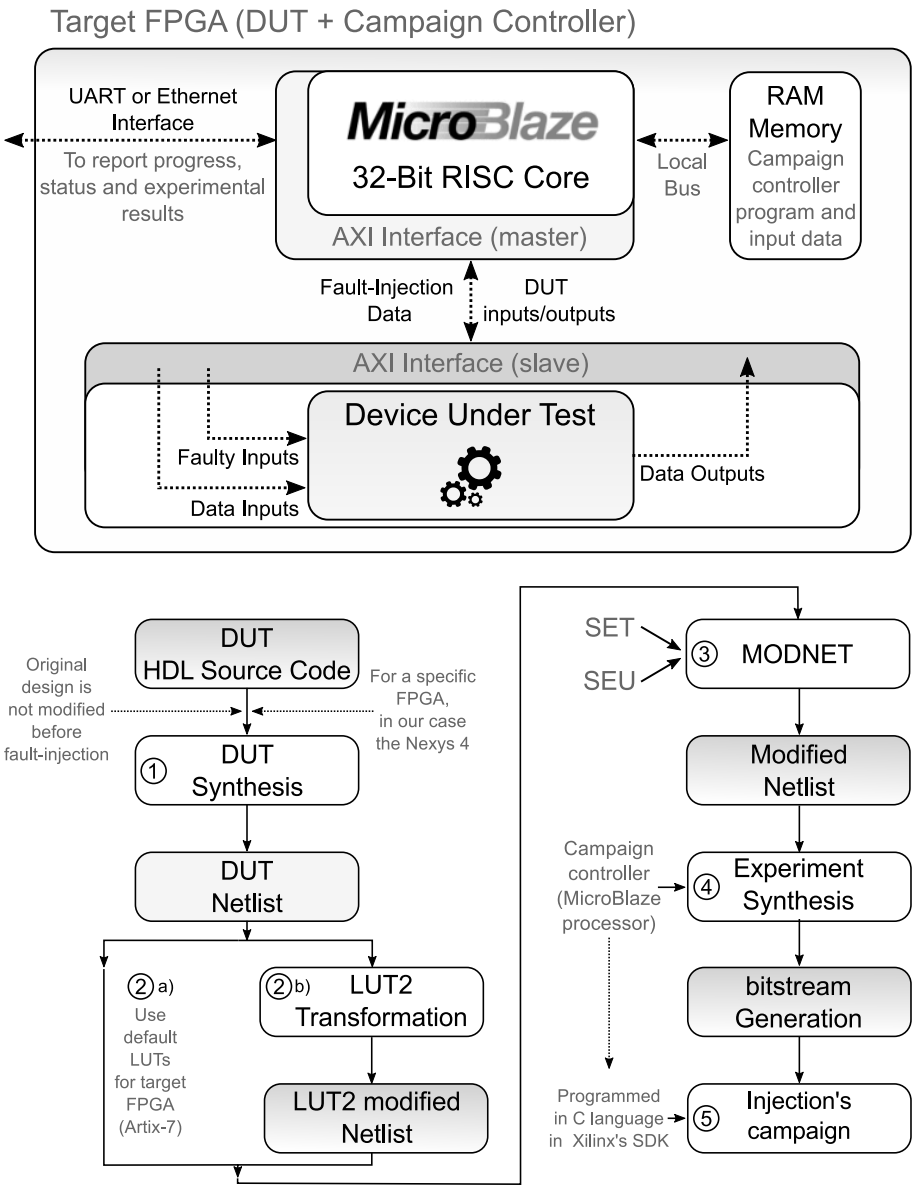

Figure A.22: NETFI-2 Framework. **Top Panel**: Architecture of the NETFI-2 framework for fault injection on FPGAs. The setup integrates the Device Under Test (DUT) and a MicroBlaze 32-bit RISC core within a single FPGA. The MicroBlaze processor manages fault injection campaigns, utilizing an AXI master-slave interface for injecting faults and communicating with the DUT. The system supports UART or Ethernet interfaces for reporting experiment results and uses onboard RAM to store campaign data. This design eliminates external controllers, streamlining the fault injection process while maintaining flexibility and scalability. **Bottom Panel**: Workflow of the NETFI-2 fault injection process. The original HDL code of the DUT undergoes synthesis, generating a netlist. The MODNET tool modifies this netlist to enable fault injection targeting SEUs and SETs. The modified netlist is synthesized with the experiment controller, generating a bitstream for the target FPGA. The experiment is executed in an environment such as ASTERICS, allowing efficient fault injection and real-time data collection. (From [Sol+17])

tially designed for long-distance and disrupted communication, to the chip scale. This involves incorporating store-carry-and-forward mechanisms and custody transfer features into chip-level communication protocols. We used theoretical reliability and availability models to compare a conventional TCP-like NoC and a DTN-like NoC. The study assumes exponential distributions for up and down times in node reliability and simulates linear topologies to highlight differences in robustness and energy consumption. The results demonstrate that DTNOCs significantly outperform traditional NoCs regarding reliability and availability. The DTN-like system exhibits four times higher mean time to failure (MTTF) and three times better steady-state availability than the TCP-like system. Additionally, the DTNOCs ensure graceful degradation, as the custody mechanism mitigates delays or faults in individual nodes, allowing the system to maintain functionality even under high failure rates.

Route Computation Unit for 3D NoCs The problem addressed in our paper in [Coe+18] is the vulnerability of Route Computation Units (RCUs) in 3D NoC to transient soft errors caused by radiation, leading to packet misrouting, deadlocks, and potential system failure. The proposed method combines a double-sampling fault detection mechanism with a rerouting strategy to correct errors before affected packets leave the router. The solution includes a custom Virtual Channel Allocator (VCA) and an additional Fault Detection Circuit (FDC) to detect illegal turns or erroneous outputs. Experimental evaluations demonstrate that the proposed architecture, implemented in a 3D-NoC with partially vertical connections, achieves robust fault correction locally, eliminating reliance on retransmissions. The design incurs a minimal area overhead of 4.1% and a power overhead of 6.8%, while maintaining competitive performance under fault injection scenarios. This approach presents a cost-effective alternative to traditional Triple Modular Redundancy (TMR) methods.

Appendix B

Analysis, Design, and Operations

This section consolidates contributions that address the fundamental questions: *How can space networks be analyzed, designed, and operated effectively?*

Space networks encompass satellite constellations, the resulting topologies defined within a contact plan, and the critical management of associated resources, such as energy and memory, required to ensure their viable and efficient operation. In light of this, this section is organized into *Constellation Analysis and Design*, **Contact Plan Design**, and **Mission Operations**.

B.1 Constellation Analysis and Design

Satellite constellations are coordinated groups of satellites designed to collectively provide global or regional coverage for communication, navigation, or Earth observation. These constellations have rapidly expanded recently, with notable examples like Starlink, OneWeb, and Planet Labs. Starlink, as shown in Figure B.1, accounts for over 60% of all operational satellites in 2024, demonstrating mega-constellations' dominance as a transformative and strategic asset in the space sector. Such networks leverage interconnected satellites to deliver seamless connectivity, underscoring their geopolitical and technological significance.

My contributions in this domain are structured into two sections: **Constellation Analysis**, which examines the operational dynamics of satellite constellations, and **Constellation Design**, which focuses on optimizing architectures, particularly for sparse configurations tailored to delay-tolerant operations with limited connectivity.

B.1.1 Constellation Analysis

Satellite constellations are complex systems that operate across vast spatial scales and extended timeframes, often spanning hours, days, weeks, or months. These dimen-

		2023		2024	
Owner	Country	Total	Share	Total	Share
SpaceX Starlink	USA	**3,395**	51%	**6,078**	63%
OneWeb Satellites	UK	502	7%	628	7%
Chinese Government	China	369	5%	369	4%
U.S. Government	USA	306	5%	306	3%
Planet Labs, Inc.	USA	**195**	3%	**240**	3%
Russian Federation	Russia	137	2%	137	1%
Spire Global Inc.	USA	127	2%	127	1%
Swarm Technologies	USA	84	1%	84	1%
Iridium Communications, Inc.	USA	75	1%	75	1%
Other	-	1,528	23%	1,528	16%
Total		6,718		9,572	

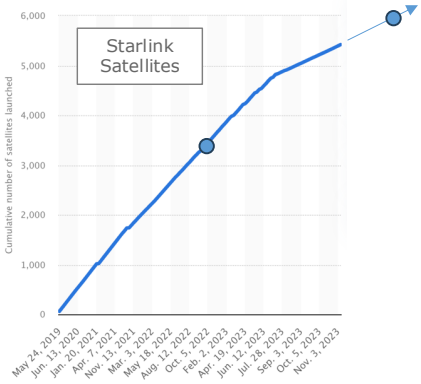

Figure B.1: Overview of mega-constellations highlighting the rapid expansion of satellite deployments from 2023 to 2024. The table shows the distribution of operational satellites by owner, with SpaceX Starlink leading in 2024, followed by OneWeb and other operators.

sions exceed typical human intuition, posing challenges to understanding and requiring innovative methods and analytical approaches to assess their dynamics effectively.

Constellation Analysis systematically evaluates satellite constellations, encompassing their spatial configurations, temporal dynamics, and operational performance. It involves assessing orbital geometries, coverage patterns, communication links, and system-level behaviors to optimize functionality, scalability, and resilience for specific applications.

This section organizes my contributions to constellation analysis into two categories: **Dense Constellations** and **Sparse Constellations**. Dense constellations feature large fleets of satellites that enable persistent, albeit dynamic, end-to-end network paths, ensuring continuous connectivity. In contrast, sparse constellations consist of smaller fleets that cannot maintain uninterrupted connectivity, relying instead on store-carry-and-forward data handling approaches aligned with the DTN paradigm introduced in Section A.

B.1.1.1 Dense Constellations

Our recent paper in [Ott+25] addresses the challenge of efficient routing in LEO and Very Low Earth Orbit (VLEO) satellite constellations, with a particular focus on ensuring polar region coverage where conventional decentralized routing algorithms underperform. VLEO refers to orbital altitudes below approximately 450 km, offering advantages like reduced communication latency and higher spatial resolution for Earth observation but posing challenges such as increased atmospheric drag, requiring innovative propulsion

and orbital maintenance technologies. Our paper conducted the first detailed sensitivity analysis of the dynamics of LEO and VLEO networks followed by state-of-the-art routing strategies, demonstrating how constellation geometry (number of satellites, orbital planes, and inclination) significantly impacts routing performance, irrespective of orbital altitude. To tackle routing inefficiencies in polar regions, we also proposed *PolarDisCo*, a novel heuristic that leverages the unique topology of polar areas to optimize inter-satellite link utilization. Although this work contributes to routing and could naturally align with Section C, its insights into VLEO dynamics, which allowed it to contribute to routing, are broader and hold greater relevance beyond routing alone. The LEO and VLEO scenarios and the range distribution across Inter-Satellite Links (ISL)s are inspired in Starlink configurations and summarized in Figure B.2. Experimental evaluations show that PolarDisCo reduces end-to-end delay by 14.8% compared to DisCoRoute in polar regions, highlighting its potential for enhancing communication efficiency in both LEO and VLEO constellations.

This analysis on VLEO networks was led by Camilla Ottaviani from Politecnico di Torino, as part of her Master's thesis (available online), which I had the privilege of co-supervising alongside Prof. Roberto Garello from the same institute. I have a strong bond with the Politecnico di Torino, as I spent several months with postdoc contracts. During my interaction with Camilla, I supported the Politecnico team in applying to the ESA tender titled *Handover, Data Routing, and Radio Resource Management for VLEO Broadband Constellations* (Reference: AO/1-11698/23/UK/ND). We won the tender, and since mid-2024, our consortium, also comprising Politecnico di Torino and Thales Alenia Italy, has been jointly addressing the critical technical challenges of VLEO broadband constellations, including handover mechanisms, routing strategies, and radio resource allocation.

B.1.1.2 Sparse Constellations

The Segmented Architecture for Earth Observation (SARE) mission exemplifies the concept of a sparse constellation. Developed by Argentinian Commission on Space Activities (Space Agency) (CONAE) as part of its New Space initiatives, the mission centered on a SA featuring a compact cluster of fewer than ten small, free-flying satellites. Unlike dense constellations with thousands of satellites providing continuous connectivity, sparse constellations like SARE inherently operate with only episodic communication opportunities between satellites and ground stations. This constraint requires innovative approaches to inter-satellite links, autonomous data sharing, and delay-tolerant operations to ensure mission success. As the primary focus of my PhD at CONAE, I analyzed and defined these critical elements, contributing to the design and operation of this cost-effective and flexible Earth observation framework.

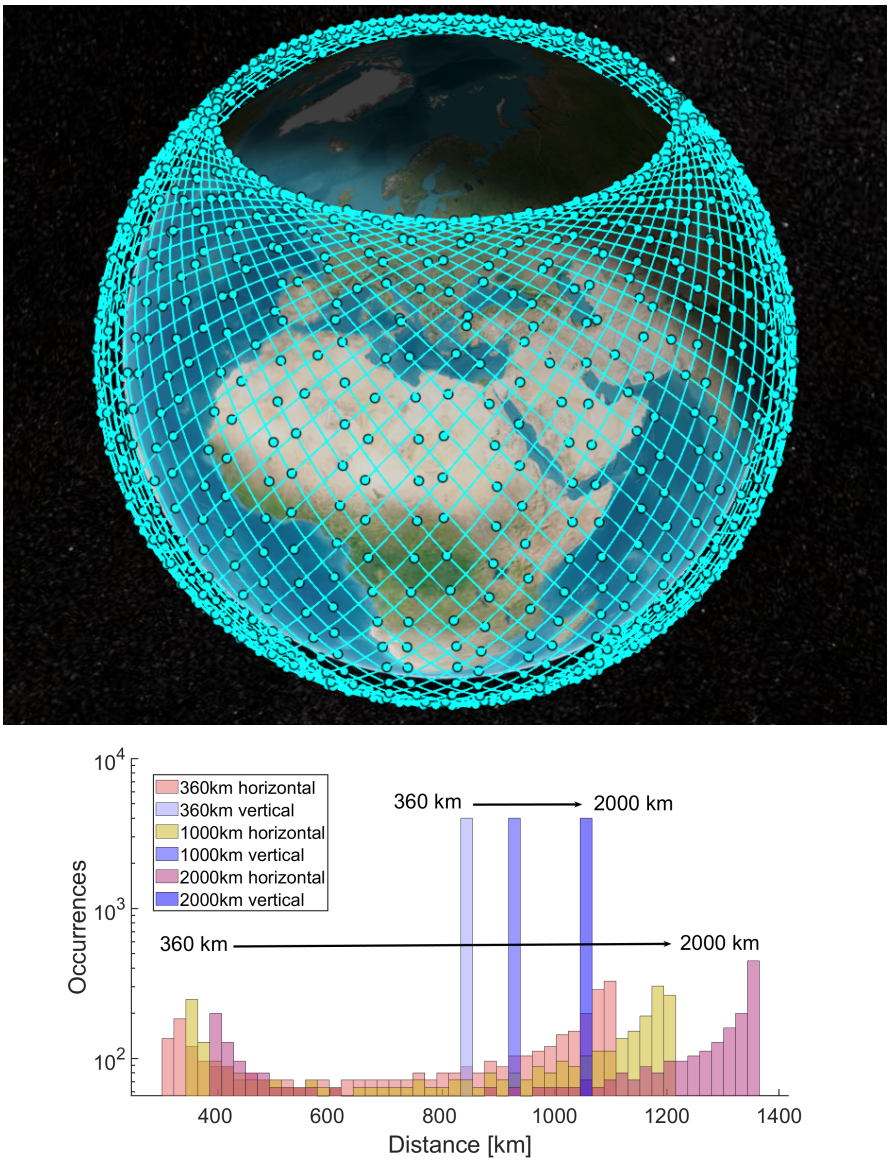

Figure B.2: Top panels depict 3D visualizations of a Starlink-like VLEO configuration with an inclination of 96.9°, 2000 satellites, 40 orbital planes, and 21 satellites per plane. The bottom panel shows the distribution of inter-satellite distances for the VLEO configuration (360 km altitude) and LEO configurations at 1000 km and 2000 km altitudes. The histograms highlight the frequency of horizontal and vertical inter-satellite distances, with VLEO exhibiting shorter distances and denser clustering than higher-altitude LEO configurations. (From [Ott+25])

Early SARE Model An early work emerging from my PhD [Fra+16] analyzed the SARE framework to identify its associated critical challenges, including managing inter-satellite communication under highly dynamic conditions, operating with limited energy resources, and ensuring system flexibility to address diverse mission requirements. The study presented a communication system for the ISL. It provided an abstract model for predicting satellite access times to ground stations and proposed a hierarchical communication design. The system design integrated DTN to manage intermittent connectivity. Simulations examined configurations with varying inter-satellite distances and evaluated access times and data throughput under multiple access control mechanisms, including TDMA and CSMA. The system achieved efficient data transfer by leveraging ISLs and DTN-based routing without requiring continuous connectivity.

Advanced SARE Model The second paper related to the SARE sparse constellation is [Fra+24], which was previously introduced in Section A.1.1 from the Earth Observation Service perspective. Besides the service, this work explores critical sparse satellite network aspects composed of just three cross-linked satellites, arranged in distinct configurations referred to as Stair (ST), along-track (AT), and cross-track (CT), as illustrated in Figure A.2. The study evaluates key topological metrics, including inter-satellite link (ISL) ranges as the swarm passes over ground stations, and assesses memory utilization under various data handling strategies such as compression and onboard processing. The primary findings of this analysis are detailed and visualized in Figure B.3.

Sparse Ring Road Networks with ISL Another paper in [FFB17b], led by myself, investigates the benefits and challenges of incorporating ISL in RRN, a low-cost architecture leveraging DTN for data transfer to remote regions. RRNs were also introduced in Section A.1.1 from the user data transport service perspective, focusing on the magazine paper in [Fel+22]. This section discusses the cross-linked sparse constellation analysis, which is the focus of this more technical paper [FFB17b]. The problem addressed focuses on improving delivery efficiency and reducing latency in RRNs, which traditionally rely on satellites acting as data mules with no direct inter-satellite communication, resulting in high delays and inefficiencies. The method simulates a case study of an RRN comprising 16 satellites distributed in a Walker constellation with ground-to-satellite and ISL capabilities. DTN protocols, particularly Contact Graph Routing (CGR), were applied to manage data flows and route optimization using a contact plan. The study evaluates metrics such as delivery ratio, latency, and network congestion through event-driven simulations in Omnet++. The key outcomes reveal that ISLs significantly enhance the delivery ratio (up to 11.7%) and reduce mean delay (up to 22.3%). However, they also introduce routing loops and congestion due to the limitations of CGR's

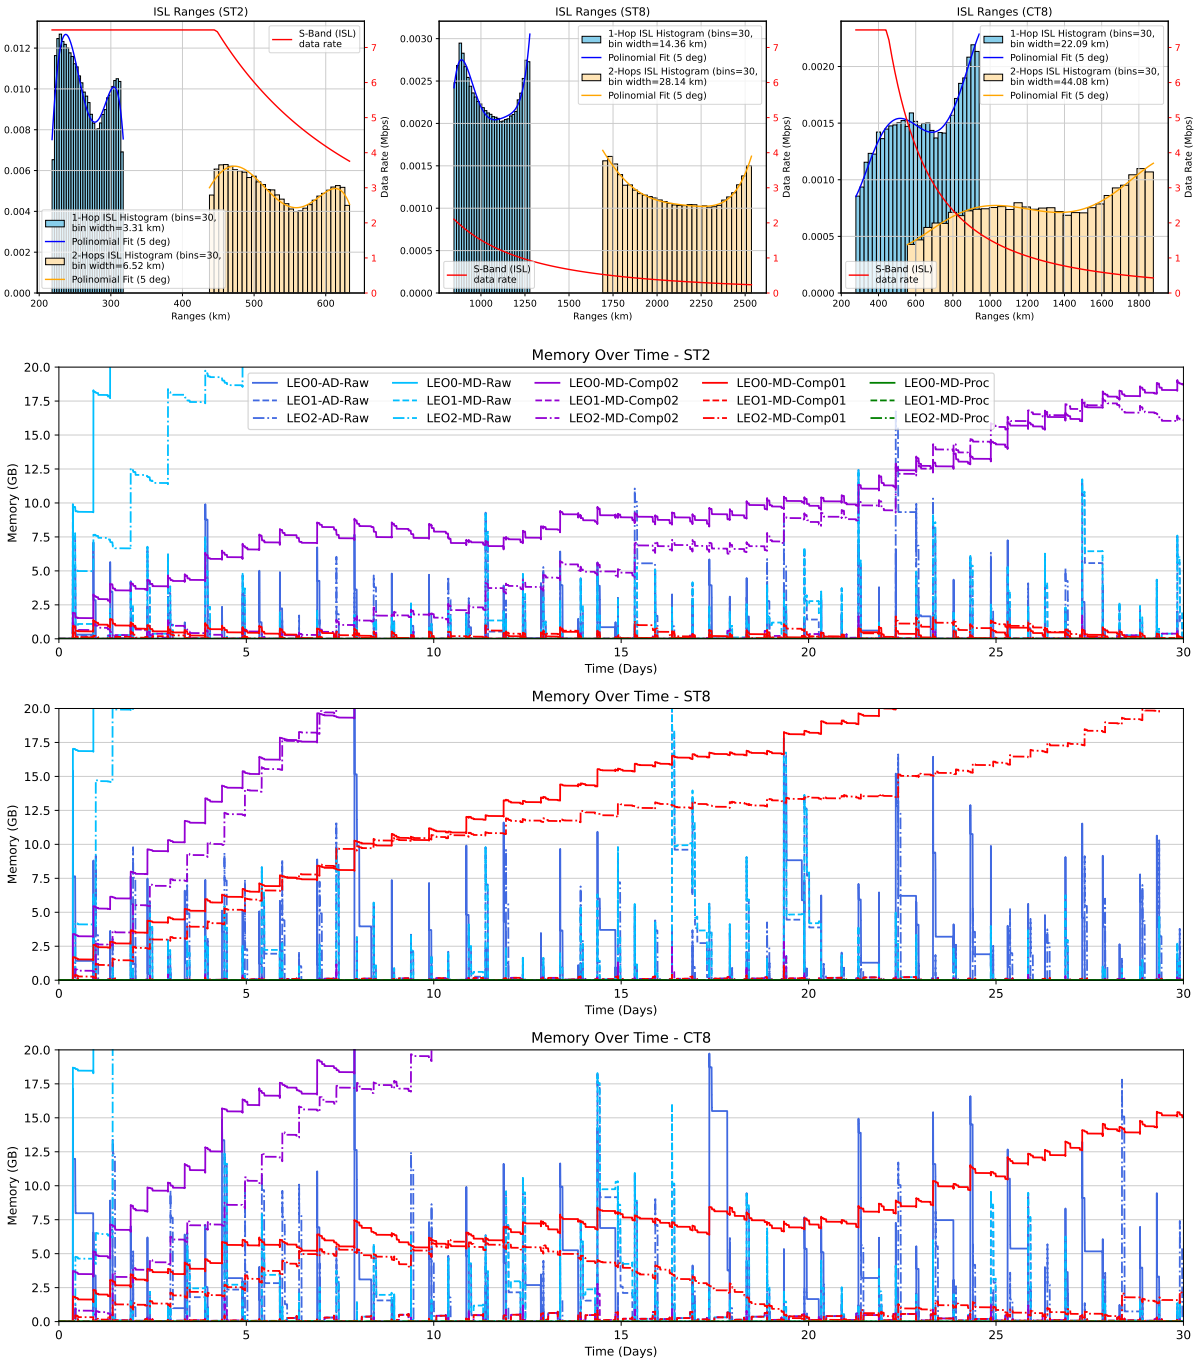

Figure B.3: SARE mission analysis. Top Panel: ISL range distributions for 1- and 2-hop distances for the selected configurations (ST2, ST8, CT8) during the satellite swarm contact with the ground stations. The ISL data rate is plotted in red for reference. Bottom Panels: Memory usage evolution over the first month for the same configurations. Raw data (RAW) exhibits the highest memory spikes during data acquisition. Compressed data (COMP) reduces these spikes by factors of 0.2 and 0.1 for two compression levels. Processed data (PROC) minimizes memory utilization to a negligible 200 bytes, keeping memory usage near zero. The results highlight the trade-offs between memory demand and data handling strategies for the three constituent satellites. (From [Fra+24])

localized congestion prediction, emphasizing the need for better congestion mitigation strategies. The results, summarized in Figure B.4, highlight the trade-offs of ISL integration in sparse satellite networks and propose congestion control as a critical area. We delve into DTN congestion contributions in Section C.2.4.

B.1.2 Constellation Design

Satellite constellation design involves optimizing orbital parameters to achieve specific mission objectives, whether for a single satellite or multi-satellite systems. This process is vital for sparse constellations, where deliberate service gaps are tolerated to accommodate delay-tolerant applications such as IoT. Sparse constellations require innovative methodologies to minimize satellite counts while ensuring operational efficiency within latency constraints.

> **Constellation Design** refers to the optimization problem of determining the optimal or near-optimal set of orbital parameters for one or more satellites within a constellation. This involves balancing coverage, latency, and resource constraints while minimizing deployment and operational costs.

This section presents my contributions to developing innovative algorithms and frameworks for designing sparse constellations tailored for IoT systems. These approaches address the challenges of delay-tolerant connectivity by leveraging heuristic optimization techniques to balance performance metrics like latency, coverage, and cost. The contributions are divided into **single-satellite missions**, where individual accessibility is maximized, and **multi-satellite missions**, where sparse constellations enable efficient IoT services on a global or regional scale.

B.1.2.1 Single-Satellite Missions

The paper in [Hen+19], led by Santiago Henn, my PhD student at CONICET in Argentina introduces the Maximum Asset Accessibility Algorithm (M3A). M3A addresses the problem of designing satellite missions that maximize access times for LEO missions constrained by limited ground station availability and high computational costs for brute force optimization. M3A is a simulated annealing-based heuristic algorithm to optimize satellite orbital parameters and ground station locations while minimizing redundant access and maximizing accessibility metrics. The method integrates with orbital propagation models to iteratively evaluate cost functions representing access time optimizations. M3A leverages probabilistic criteria to explore the search space efficiently, avoiding local optima while incorporating constraints such as ISL capabilities and ground sta-

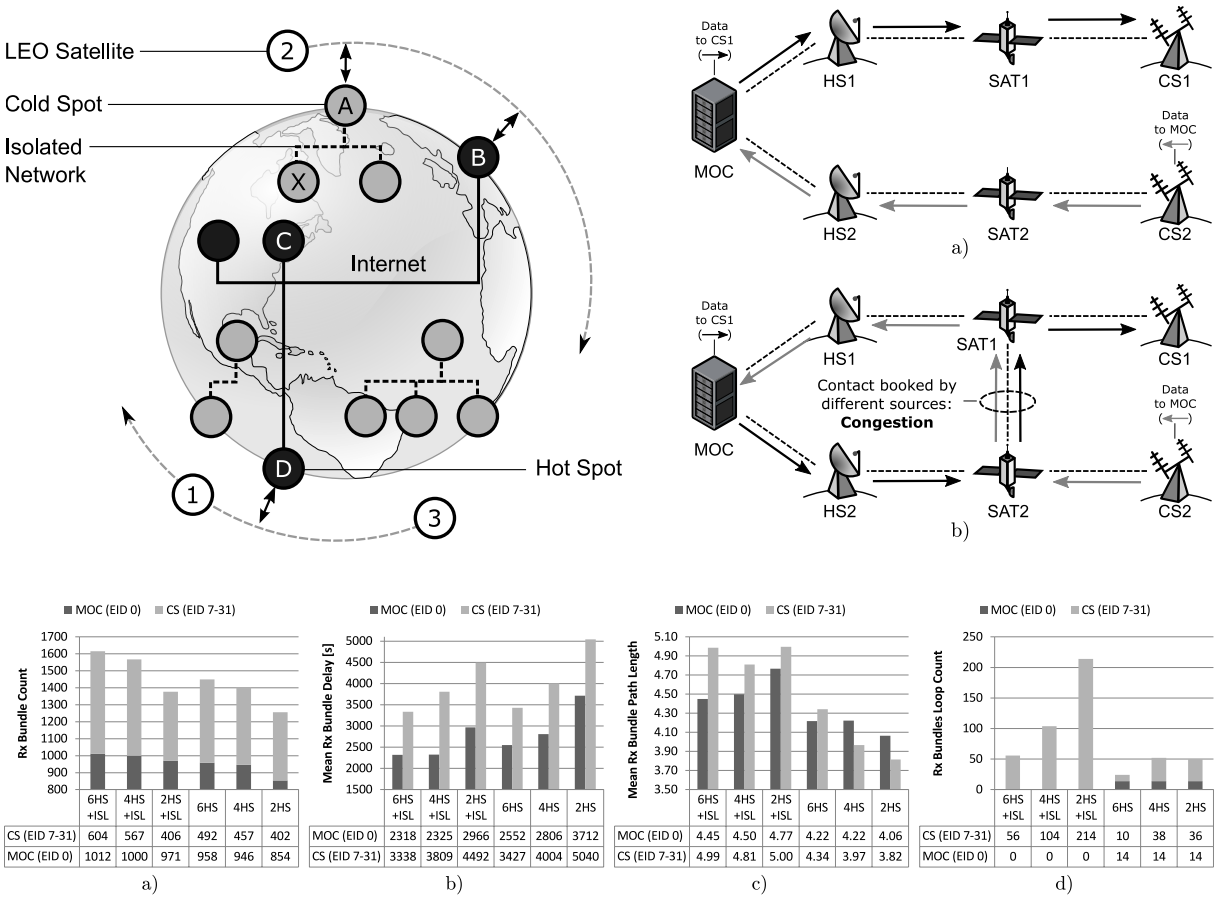

Figure B.4: Ring Road Network (RRN) analysis. Top Left: The RRN example consists of a set of LEO satellites providing connectivity between isolated areas (cold spots, CS) and centralized hubs (hot spots, HS). Cold spots are geographically remote or underserved regions lacking direct access to the Internet, while hot spots are connected nodes with direct Internet access (e.g., ground stations). LEO satellites act as data mules collecting, storing, and delivering data to and from cold spots and hot spots. Top Right: Congestion problem due to routing loops introduced by ISL. (a) Without ISLs, data flows directly from satellites (SAT1 and SAT2) to their assigned ground stations (HS1 and HS2). (b) With ISLs, contact bookings from multiple sources can cause congestion in satellite links, resulting in routing loops and delays. Bottom: Performance results for 2, 4, and 6 hot spots (HS) with and without ISL, focusing on key metrics: (a) Rx-bundle Count (the number of received Bundles), (b) Mean Bundle Delay (seconds of in-transit time of received bundles), (c) Mean Bundle Path Length (number of hops), and Bundle Loop Count (the number of routing loops encountered). The results confirm that while ISLs significantly enhance the throughput of RRNs, they also introduce routing challenges. The evaluated CGR exhibits critical routing loop flaws that must be resolved to deploy cross-linked RRNs effectively. (From [FFB17b])

tion placements. A case study simulating a LEO relay mission with one primary mission satellite, three relay satellites, and two ground stations demonstrates the algorithm's effectiveness. Our study compares M3A's results with brute force optimization to validate its performance. The key outcomes reveal that M3A achieves near-optimal solutions with significantly reduced computational time—23 times faster than brute force—while yielding only a 0.43% deviation from the optimal cost function. The relay network enabled an additional 1527 minutes of real-time access, substantially improving ground station-only configurations. The study underscores M3A's capability to streamline the design of space missions, making it particularly valuable for small agencies and commercial projects with limited budgets.

B.1.2.2 Multi-Satellite Missions

Sparse constellation design was not appropriately addressed in satellite network architecture, although crucial for delay-tolerant IoT applications. Unlike dense constellations, sparse constellations intentionally allow periodic service gaps between satellite passes, leveraging the delay tolerance inherent to many IoT systems. Multi-satellite sparse constellation design approaches are crucial to minimizing the number of satellites while maintaining service continuity within acceptable latency bounds. My contributions allowed me to fill this gap in research as follows.

Global IoT Constellation Design The first paper in [Fra+20a], led by myself, addressed the challenge of designing satellite constellations for IoT systems using the LoRaWAN protocol, where continuous connectivity is not required, but latency constraints must be respected. The goal is to minimize the constellation size while ensuring compliance with LoRaWAN Class-B requirements for maximum beacon-less operation time. In this work, I first introduced the concept of sparse constellations, leveraging the inherent delay tolerance of LoRaWAN Class-B devices. To ensure global or regional coverage with minimal infrastructure, I propose a heuristic-based complexity-progressive gradient descent optimization algorithm that determines the orbital parameters, including the number of satellites, planes, and inclinations. The method uses parametric case studies to identify configurations that balance latency and deployment cost. Results demonstrate that a global LoRaWAN service can be achieved with only nine satellites, providing a 2-hour gap coverage compliant with the protocol.

Regional IoT Constellation Design The subsequent paper in [Cap+22], led by Gabriel Maiolini Capez, my PhD candidate at the Politecnico di Torino, explores sparse constellation design for networks also employing NB-IoT besides the LoRaWAN protocol. This

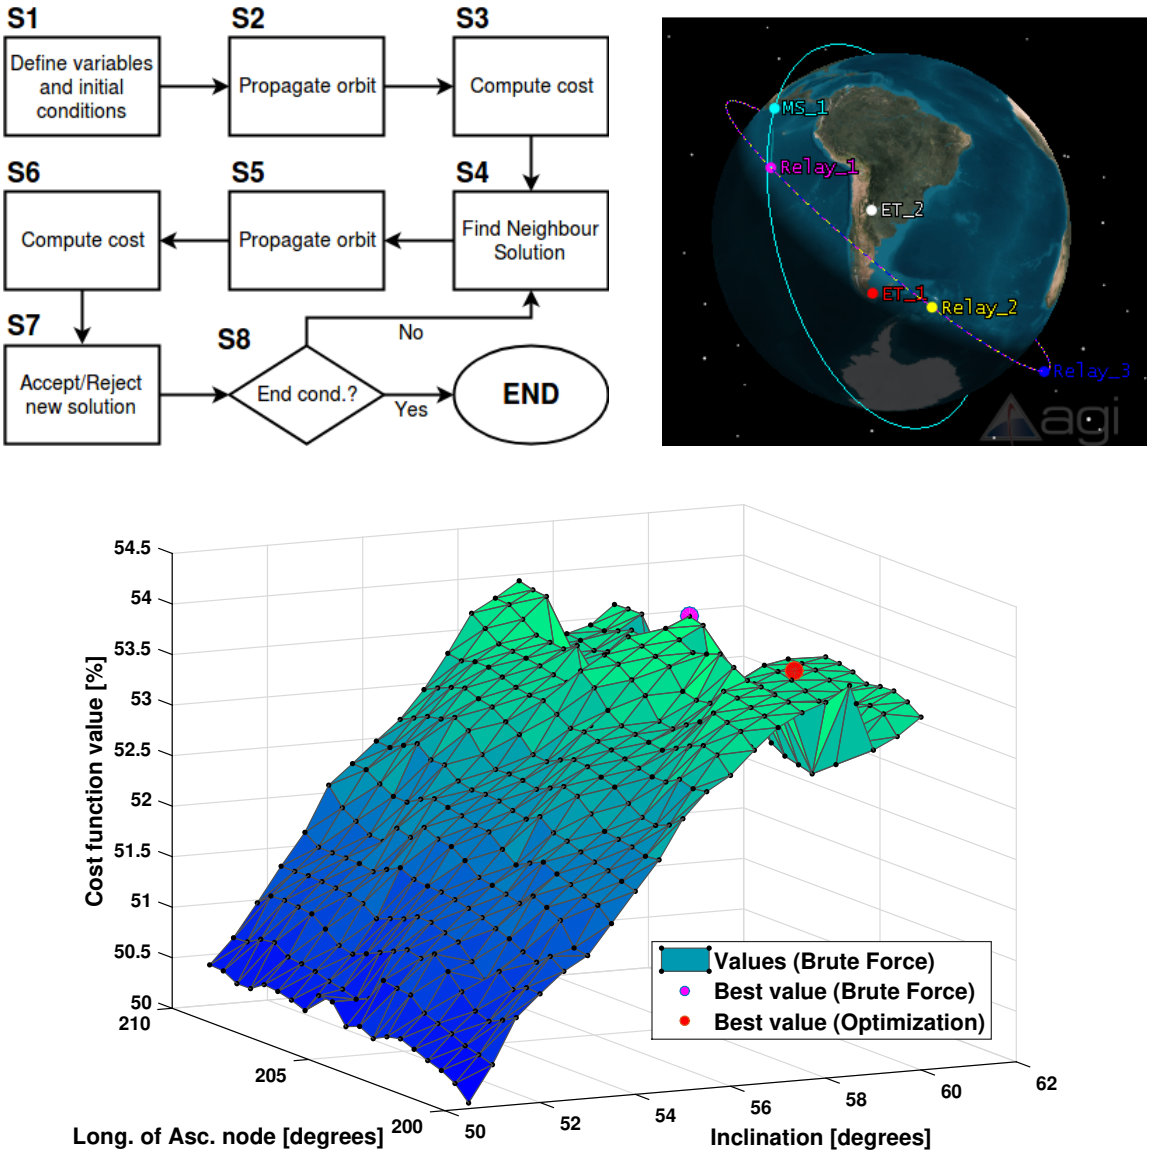

Figure B.5: Maximum Asset Accessibility Algorithm (M3A). Top Left: The M3A algorithm flowchart outlines the simulated annealing-based heuristic optimization process. Steps include defining variables and initial conditions (S1), orbit propagation (S2), cost computation based on accessibility metrics (S3), finding neighboring solutions (S4), and evaluating the acceptance criteria for convergence (S7-S8). This iterative process continues until the stopping conditions are met. Top Right: The resulting orbital configuration shows the optimized positions of the main satellite (MS_1) and three computed relay orbits (RELAY_1, RELAY_2, and RELAY_3). These orbits ensure extended accessibility to the ground station network, represented as ET_2. Bottom: The cost function across the optimization space. The function comprises metrics such as total access time, relay efficiency, and ground station utilization. The global optimum (blue dot) determined by brute force and the best value (red dot) found by the simulated annealing M3A approach demonstrates that M3A achieves a near-optimal solution. (From [Hen+19])

study also focuses on optimizing constellation design for global and regional coverage while addressing constraints like maximum revisit times and satellite positioning. The gradient descent approach from [Fra+20a] was adapted to consider the unique characteristics of NB-IoT and LoRaWAN protocols, such as maximum beacon-less operation time and power-saving modes, to achieve quasi-optimal satellite configurations. Both global and regional scenarios are evaluated, including innovative approaches like hexagonal meshing for regional coverage. The LoRaWAN beaconless mode and NB-IoT power-saving modes that define the gap constraints are illustrated in Figure B.6. The flow of the adapted gradient-descent algorithm is also plotted in the figure. Following the first paper, sparse constellations were shown to provide global IoT coverage with significantly fewer satellites than dense constellations. For instance, NB-IoT services can achieve global coverage with just six satellites (a 3-hour gap), while LoRaWAN-based systems require nine satellites (a 2-hour gap). Figure B.7 provides a comprehensive summary of the study's outcomes, showcasing the trade-offs between constellation size and algorithmic cost regarding computational time and memory usage. Regional configurations with as few as three satellites demonstrated adequate coverage for Africa. These findings highlight the cost-efficiency, regional flexibility, and scalability of sparse constellations for IoT applications.

B.2 Contact Plan Design

The constellation analysis and design contributions in Section B.1 primarily focus on satellite constellations' physical, geometrical, and geographical characteristics. In contrast, Contact Plan Design (CPD) focuses on the immediate network abstraction derived from these constellations: the topological model. This model identifies which nodes can communicate with one another and under what conditions, considering architectural, energy, interference, or other constraints. These constraints necessitate further topological analysis to realize an effective functional and efficient space network.

This section of the HDR holds particular significance, as I was the one to coin the term "Contact Plan Design" (CPD) in 2015, as detailed in [FF15a] and further revised in the following Sections B.2.1 and B.2.2. During my PhD involvement with the SARE mission at the Argentinian Space Agency, I identified a critical challenge in distributed network operations related to contact planning. This realization allowed me to abstract and generalize the problem into a formal CPD framework, enabling the strategies and solutions developed to extend beyond SARE and apply to a wide array of space networks. This foundational work has catalyzed a substantial body of research, leading to numerous publications by my students and researchers worldwide. It is profoundly fulfilling to witness how CPD approaches are being adapted and implemented in space

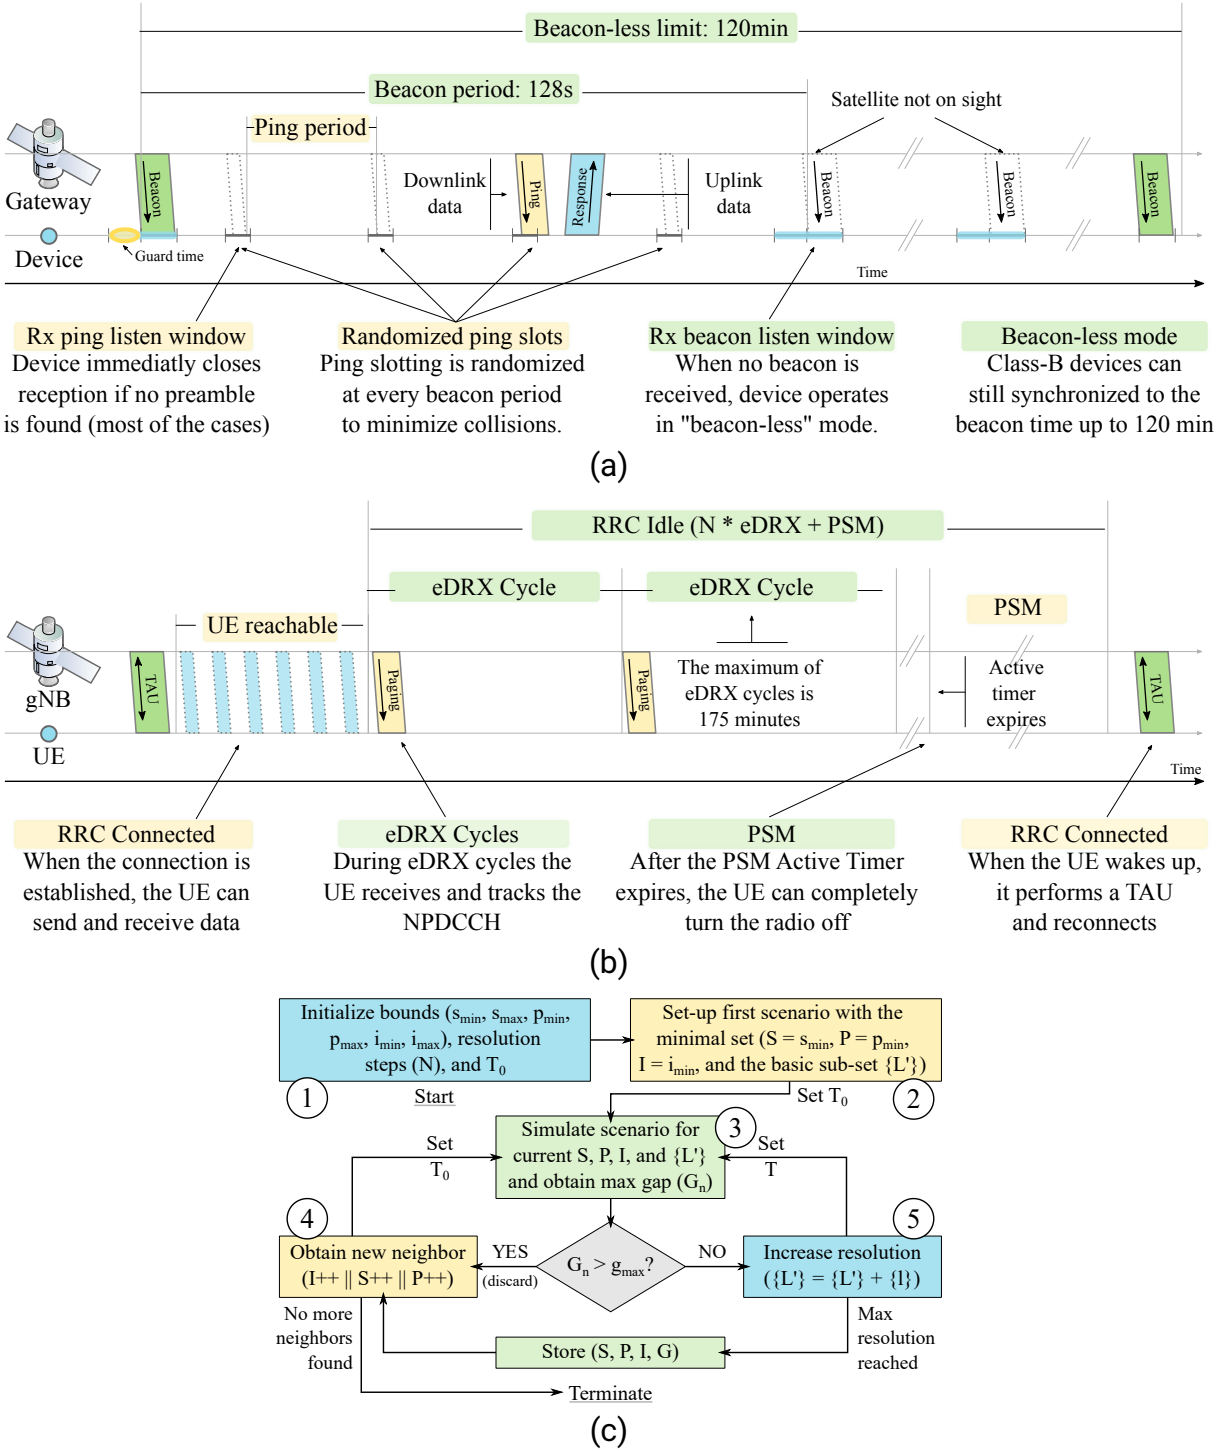

Figure B.6: Sparse IoT constellation constraints and design. (a) **LoRaWAN Beacon-Less Mode**: Class-B devices operate without beacon synchronization, listening during randomized ping slots to minimize collisions and allowing up to 120 minutes of operation without a gateway in sight. (b) **NB-IoT eDRX and PSM Modes**: These modes optimize energy consumption by enabling periodic downlink monitoring during eDRX cycles and powering off the device radio entirely in PSM until the next access update. (c) **Gradient Descent Algorithm**: An iterative approach for sparse DtS-IoT constellation design, refining satellite count and orbital parameters to minimize communication gaps while balancing performance metrics. (From [Cap+22])

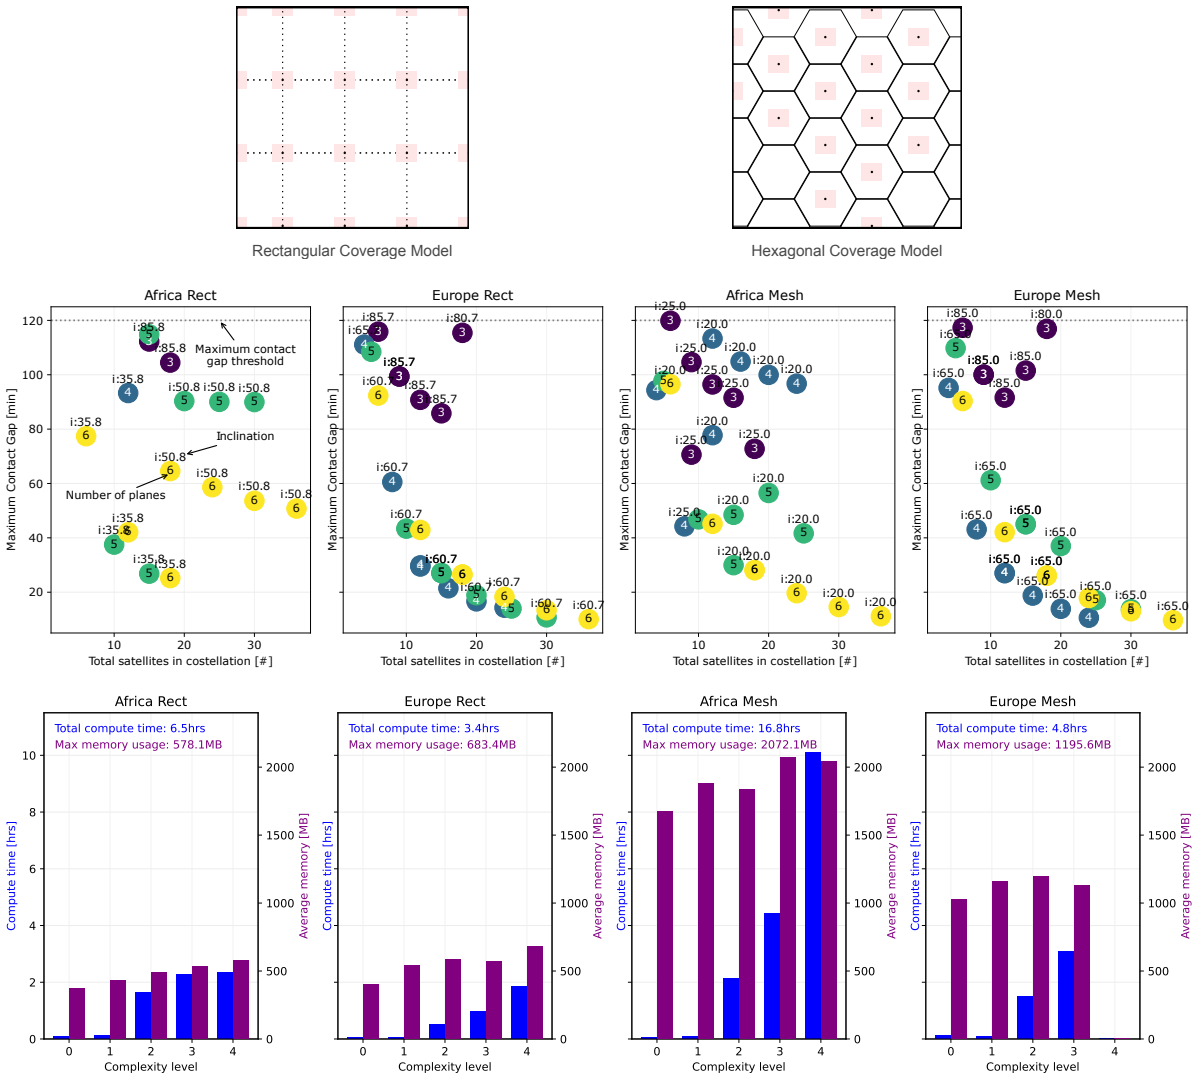

Figure B.7: Performance evaluation of the sparse DtS-IoT constellation design algorithm for Europe and Africa regions under a maximum contact gap constraint of $g_{max} = 120$ min. **Top Panel**: Rectangular and hexagonal coverage models studied in the paper. Center row: The maximum contact gap is plotted against the number of satellites for boundary configurations of rectangular (left two panels) and hexagonal mesh (right two panels) models studied in [Cap+22]. Each data point represents a specific satellite constellation configuration, where the number inside the circle indicates the number of orbital planes, and the value next to "i:" specifies the inclination in degrees. The results show that smaller satellite counts approach the maximum gap threshold tolerated by LoRaWAN and NB-IoT protocols. Hexagonal mesh configurations generally yield better performance, achieving smaller gaps with fewer satellites than rectangular bounds. **Bottom Panel**: Resource consumption metrics, including compute time (purple bars) and memory usage (blue bars), are presented for increasing levels of algorithmic complexity. Hexagonal mesh configurations demand higher computational and memory resources, particularly for Africa, than rectangular configurations. These results highlight the trade-offs between computational cost and improved constellation design quality. (From [Cap+22])

networks across multiple nations, underscoring the global impact of this contribution.

One of my key future actions will be to develop a comprehensive survey that revises and consolidates all CPD methods created over the past decade. This survey will serve as a critical resource for the space networking community, providing a thorough overview of advancements, methodologies, and applications while identifying gaps and opportunities for further research in this essential field.

B.2.1 Problem Definition

The Contact Plan Design (CPD) process was first formally introduced during my early PhD research [FF15a], where its significance for managing communication in space networks was established.

Contacts are the key constituent element of the CPD process. A *contact* is formally defined as a temporary link between two nodes, such as satellites or ground stations, satisfying specific physical and operational conditions. This definition underscores the transient and resource-sensitive nature of space network communications. A *contact plan* is a curated set of scheduled contacts within a specified period, defining the topology of the forthcoming space network and serving as a foundation for reasoning, optimization, and control of network operations. With this in mind, we can define CPD as follows.

> **Contact Plan Design** is the systematic process of selecting and scheduling communication opportunities, known as contacts. CPD transforms raw connectivity data from orbital mechanics into an optimized contact plan by considering constraints like energy budgets, antenna pointing capabilities, interference avoidance, and mission-specific requirements (see Figure B.8).

The outcome of CPD is a structured contact plan that enables efficient data transfer across dynamically evolving network topologies. This process balances key performance metrics such as throughput, latency, and fairness while adhering to stringent resource limitations. CPD is critical in addressing the complexities of intermittent and resource-constrained communication in space networks by providing a predictive and methodical framework for network management.

Given the diversity of constraints and the myriad optimization possibilities, a wide range of CPD methodologies has emerged over the past decade. The methods I contributed are categorized and discussed in Section B.2.2.

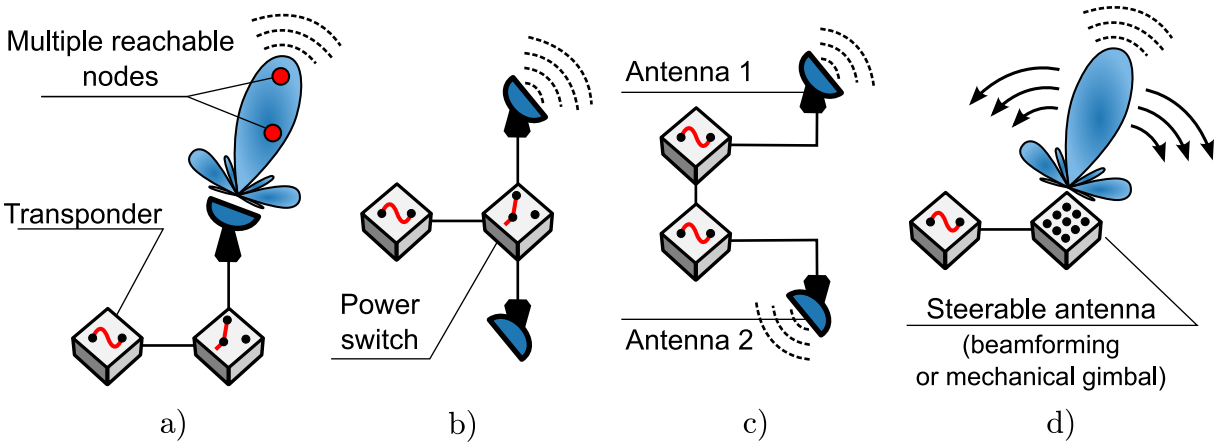

Figure B.8: Key constraints that govern the CPD process: (a) **Reachable Node Determination**: While a line of sight may exist between satellites and ground stations, most satellite communication links are designed to be point-to-point rather than multiple-access due to the limitations at such distances. Consequently, the CPD process must determine which reachable or illuminated nodes the link will be effectively established with, ensuring efficient utilization of the available connectivity. (b) **Antenna Determination**: Many spacecraft, such as those in the SARE mission, are equipped with multiple antennas connected to a single transponder via a power switch. This configuration enables link establishment in multiple directions without relying on body-pointing (where the satellite physically rotates to adjust its orientation). The CPD process is essential for selecting the appropriate antenna and ensuring the power switch is pre-configured to support the desired link. (c) **Transponder Determination**: More capable missions could embark on multiple transponders, each connected to its corresponding antenna, which solves the antenna determination problem. However, in many cases, other constraints arise. Notably, the energy limitation that, in some cases, forbids the simultaneous utilization of multiple transponders. Other reasons, such as interference or electromagnetic compatibility, might force a transponder decision. (d) **Pointing Determination**: For spacecraft equipped with mechanically gimbaled antennas (motorized systems that rotate high-gain antennas) or electronically steerable antennas (using phased-array technology to direct beams), precise pointing decisions must be made. These decisions involve selecting the optimal neighbor for communication and configuring the antenna direction accordingly, a process that is inherently part of CPD. (a, b, and c) **Determination Frequency**: In all cases, these selections occur across multiple networked spacecraft within a space network, making the process inherently interdependent across the entire system. Given the dynamic nature of space networks, where topologies are continuously changing, the CPD process must be executed frequently, often with every topology update. This underscores the need for algorithmically efficient solutions to handle the complexity and ensure timely and effective decision-making. (From [FF15a])

B.2.2 Contact Plan Design Methods

CPD methods are essential for transforming raw satellite connectivity data into structured and optimized communication schedules. This section categorizes and discusses various CPD methodologies developed to address distinct operational challenges in satellite networks. These methods consider **topology** information, **routing** behavior, **traffic** distribution, **energy** constraints, and **positioning** data to optimize space network performance. All of these efforts were led by myself during my PhD thesis under the supervision of Prof. Jorge M. Finochietto and Prof. Pablo A. Ferreyra.

B.2.2.1 Topology-Awareness

Topology-aware CPD methods focus on ensuring fairness and efficiency in network connectivity, relying solely on satellite configurations and their interconnectivity. These methods are the most straightforward CPD approaches, requiring no additional input beyond the connectivity graph of the envisioned space network. They also mark my initial contributions to the CPD field.

> I owe special acknowledgment to Prof. Jorge M. Finochietto, one of my PhD supervisors, whose guidance emphasized tackling more minor, manageable problems as a foundation for addressing more complex ones. This approach sharply contrasted with my youthful ambition to solve entire problems in one go, considering all possible inputs—a strategy that inevitably led to confusion and frustration. This iterative approach remains a cornerstone of how I mentor my students today.

The paper in [FMF13] addresses the challenge of designing fair CPD in DTN, where connectivity is intermittent but predictable, such as in satellite constellations. The core problem is allocating contacts in a time-evolving topology that maximizes fairness across nodes while maintaining network capacity. I proposed a multi-objective optimization model formalized as a Mixed-Integer Linear Programming (MILP) problem to balance fairness and capacity, considering constraints like limited interfaces per node. Recognizing the computational intractability of the MILP model for large-scale systems, I developed an efficient algorithmic solution based on the *Blossom algorithm*, adapted to handle non-perfect matchings and maximize fairness over time by incorporating a weighted mechanism for underutilized contacts. The Blossom-based method was coined Fair Contact Plan Design (FCP) and became the baseline for other CPD works such as [Yan+20] addressed later in Section B.2.2.5. The evaluation shows that the proposed algorithm achieves near-optimal fairness with polynomial complexity while

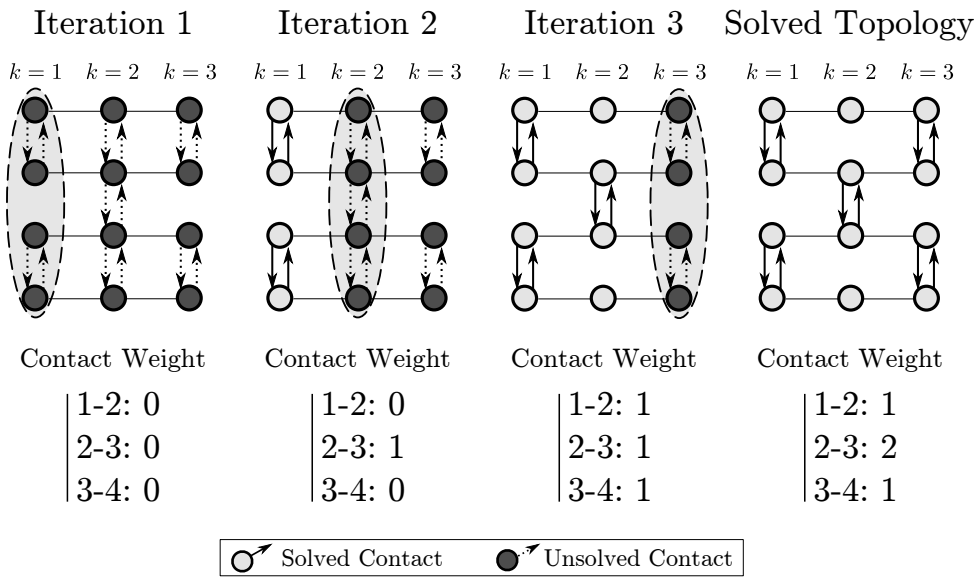

Figure B.9: Fair Contact Plan (FCP). The figure illustrates the iterative process of the FCP algorithm, which progressively designs the connectivity graph of a dynamic topology. At each iteration, the algorithm evaluates the contact weights and filters out contacts that penalize fairness, ensuring a balanced and efficient network. The Blossom algorithm is employed at each step to optimize the maximum matching within the connectivity graph. (From [FMF14a])

maintaining system capacity.

A journal extension of this paper followed [FMF14a]. The journal extends the FCP framework's applicability beyond its initial focus, suggesting its potential for other space network configurations and traffic patterns. The extended work includes several practical case studies of a distributed orbiting sensor network. It applies the developed methods to real-world scenarios, assessing the impact of FCP on routed traffic, such as average bundle age, under realistic conditions simulated with the OMNeT++ framework[1].

B.2.2.2 Route-Awareness

The subsequent paper in [FF15b] builds upon the concept of fair contact planning by addressing the problem of designing routing-aware contact plans. While focusing on fairness in link assignment among nodes, this work extends the approach to minimize all-to-all end-to-end multi-hop route delays, integrating routing efficiency into the contact selection process. I coined the method as Route-Aware Contact Plan Design (RACP). To tackle this multi-objective optimization problem, the paper proposes a two-stage method. First, the Fair Contact Plan Design (FCP) algorithm presented in the previ-

[1]CPD public repository: https://github.com/juanfraire/dtn-topology-solver

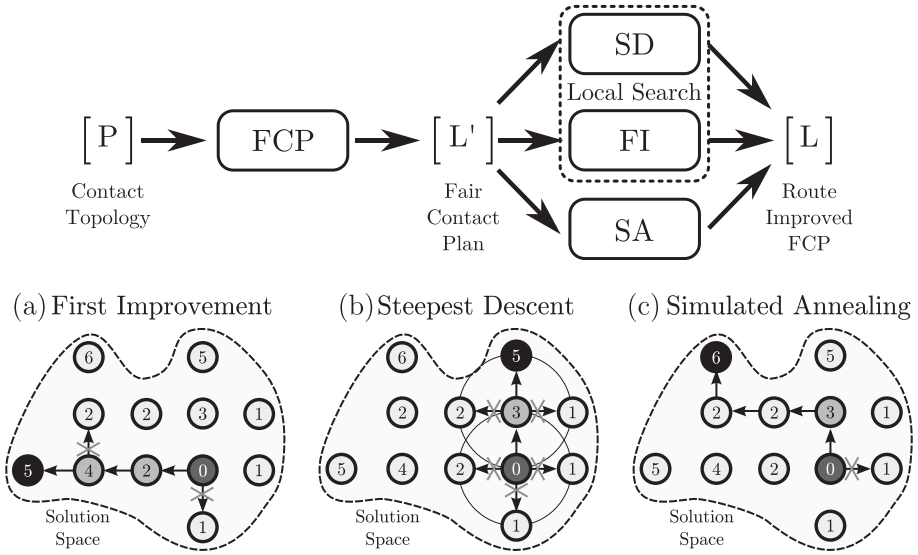

Figure B.10: Route-Aware Contact Plan Design (RACP) Process. **Top Panel**: The overall flow of route-aware CPD begins with the Fair Contact Plan (FCP) generation, which ensures fairness in link assignments. Subsequently, heuristic methods such as First Improvement (FI), Steepest Descent (SD), and Simulated Annealing (SA) iteratively refine the initial contact plan to optimize routing metrics, producing a route-improved contact plan. **Bottom Panel**: Visualization of how different heuristics explore the solution space of the CPD problem. (a) FI: Incrementally navigates the solution space by accepting the first improvement found. (b) SD: Explores all possible moves in the solution space and selects the most optimal one at each step. (c) SA: Employs a probabilistic approach to escape local optima by accepting non-optimal solutions occasionally, enabling a broader exploration of the solution space.

ous section is applied to distribute contacts equitably over time, leveraging the Blossom algorithm to achieve a balanced matching at each topology state. Second, local search techniques, including First Improvement (FI), Steepest Descent (SD), and Simulated Annealing (SA), refine the contact plan by incorporating routing metrics like delay and unrouted time while maintaining fairness. Figure B.10 presents an overview of the route-aware CPD methods. Metrics include Fairness Index (JI) to measure the equitable distribution of link opportunities, Maximum Average Delay (MD) to assess the worst-case end-to-end multi-hop route delays, and Unrouted Time (UT) to quantify periods of network disconnection. The evaluation reveals that while FI and SD efficiently improve FCP solutions, SA outperforms these methods by achieving Pareto-optimal solutions that balance fairness, routing efficiency, and computational complexity. The results underscore the importance of routing-aware CPD in enabling effective communication in constrained DTN scenarios, with applications in real-world satellite networks such as Earth observation constellations.

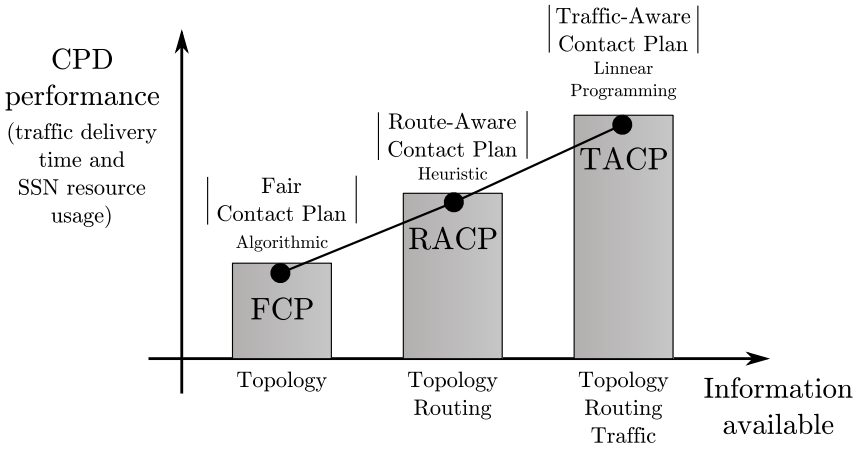

Figure B.11: Conceptual evolution of Contact Plan Design (CPD) methods based on available information and performance metrics. The Fair Contact Plan Design (FCP) utilizes topological data only. Route-Aware Contact Plan Design (RACP) incorporates both topology and routing information. Traffic-Aware Contact Plan (TACP) integrates topology, routing, and traffic demand information, achieving the highest CPD performance regarding traffic delivery time and resource efficiency for Space Sensor Networks (SSNs). (From [FMF16])

B.2.2.3 Traffic-Awareness

Having explored topology-aware CPD methods, which focus solely on the structural properties of the contact plan, and routing-aware approaches, which incorporate end-to-end route metrics to refine the contact plan, the next logical progression is traffic-aware methods. These methods advance the CPD framework by considering the actual traffic demands expected during the contact plan period, enabling more informed decisions on which contacts to prioritize or filter. Traffic-aware approaches thus require additional input information, specifically the anticipated traffic flows, to optimize the network's performance under realistic operating conditions. These aspects are illustrated in Figure B.11.

TACP My contribution in [FMF16] addresses optimizing contact plans while considering network constraints like storage, transmission capacities, and single-link limitations. The proposed solution, Traffic-Aware Contact Plan (TACP), is formulated as a MILP problem. TACP integrates predictable topology and traffic demands to minimize delivery delays and optimize resource usage at intermediate nodes. A case study demonstrates the superiority of TACP over prior methods such as FCP and RACP (both discussed in the previous paragraphs), achieving a 42% improvement in resource efficiency and a 10% reduction in delivery time for the evaluated scenario. A comprehensive analysis of the results presented in this paper is illustrated in Figure B.12. These results highlight TACP's effectiveness in enabling efficient and scalable data delivery for Earth

observation missions.

TACP-EA One of the primary challenges of TACP lies in the significant computational effort required to solve its MILP formulation due to its NP-complexity, particularly when contrasted with the more efficient Blossom-based algorithmic approach employed by FCP. I focused on developing heuristic alternatives for TACP to address this limitation during my PhD. The initial findings from this effort were published in [Fra+15b], with the final and refined results appearing in [Fra+17c]. In these papers, I proposed Traffic-Aware Contact Plan based on Evolutionary Algorithms (TACP-EA), a TACP framework based on Genetic Algorithms (GA), as a heuristic alternative to the MILP-based TACP formulation. The TACP-EA utilizes a population-based approach inspired by genetic evolution, incorporating repair strategies for infeasible solutions and optimization criteria based on traffic delivery times. Key outcomes demonstrate that TACP-EA significantly reduces computational time while producing implementable contact plans with near-optimal performance. More details on the model, encoding, and hyperparameter tuning performed in this paper are provided in Figure B.13. These results prove that the proposed algorithm performs well under bounded time constraints, making it suitable for real-world satellite missions that require timely contact plan updates.

B.2.2.4 Battery-Awareness

In topology-, routing-, and traffic-aware CPD methods, constraints were predominantly defined by the maximum number of simultaneous contacts a given node could sustain at any given moment. Addressing these constraints effectively marked the culmination of the primary objectives of my PhD research.

After completing my PhD, I had the privilege of collaborating with the Dependable Systems and Software department at Saarland University, Germany, led by Prof. Holger Hermanns. Their expertise in battery modeling through model-checking methodologies, particularly within the context of projects like SENSATION, provided a unique perspective and set the stage for one of the most significant international cooperations of my career. This collaboration extended my research into the domain of battery-aware CPD, integrating energy consumption models into the design of sustainable space network solutions.

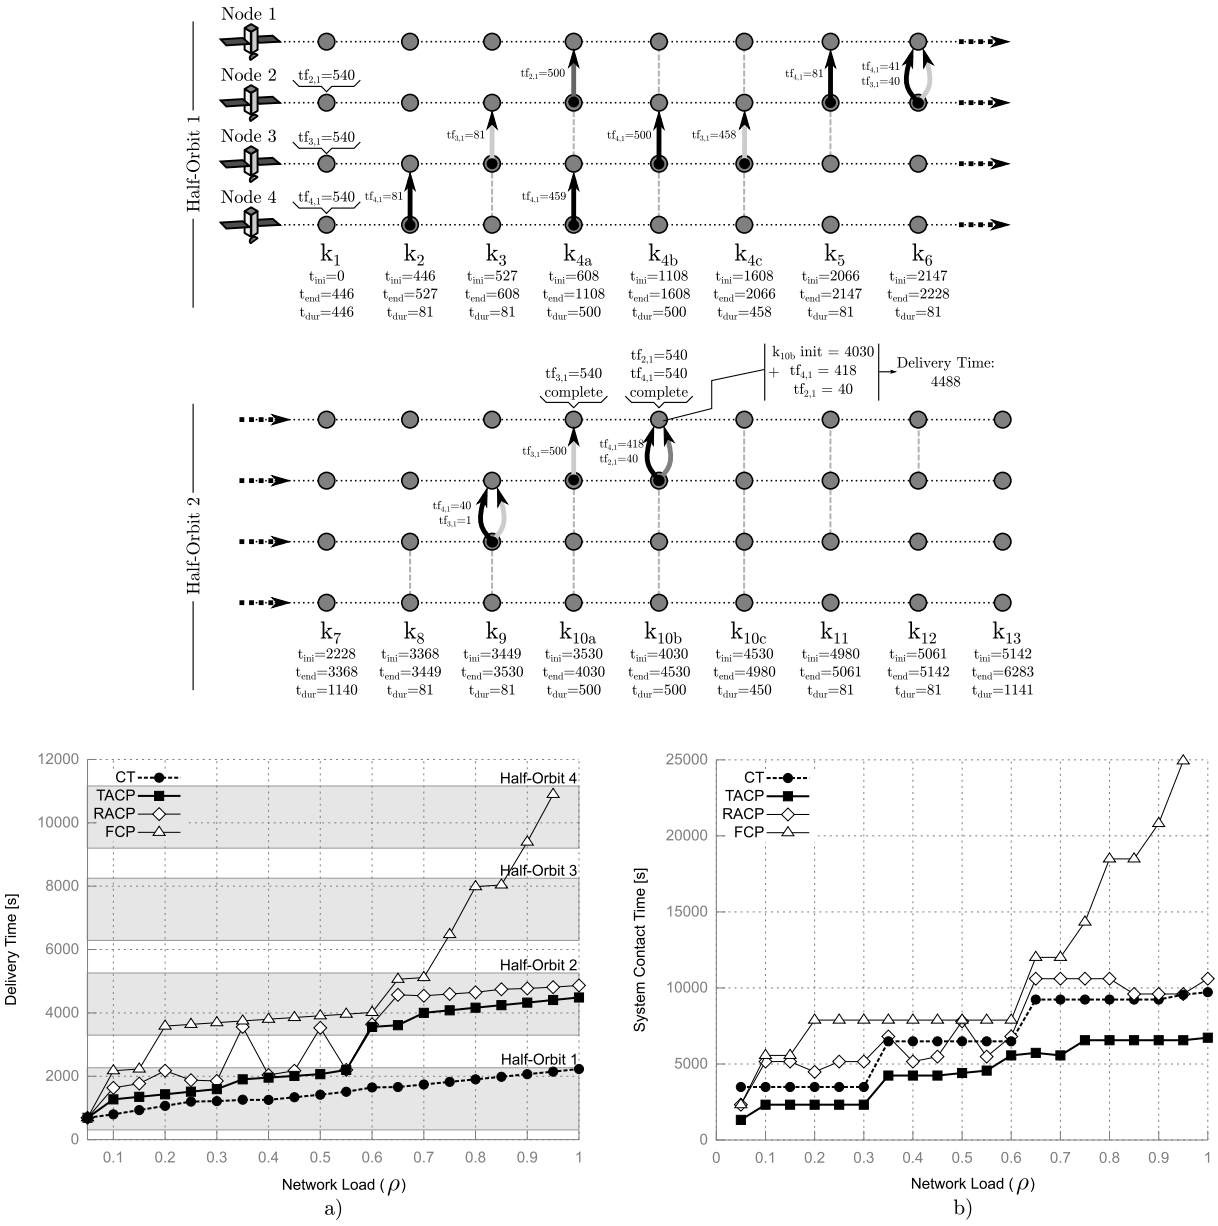

Figure B.12: Traffic-Aware Contact Plan (TACP) performance evaluation. **Top Panel**: An example contact plan output from TACP, showing the allocation of multiple traffic flows ($tf_{src,1}$) with Node 1 as the destination. Contacts in the plan are assigned based on traffic demands and available network capacity, optimizing data delivery times. **Bottom Panel**: Comparative analysis of delivery time (a) and accumulated system contact time (b) for increasing network loads across different CPD methods. TACP systematically achieves the lowest delivery latency by leveraging traffic awareness, outperforming Route-Aware Contact Plan Design (RACP) and Fair Contact Plan Design (FCP). Physical contact time (CT), representing the theoretical minimum delivery time without CPD constraints, is provided as a baseline. Notably, while FCP maximizes topological fairness, it results in the highest accumulated contact time without achieving the delivery time improvements enabled by traffic-aware planning. (From [FMF16])

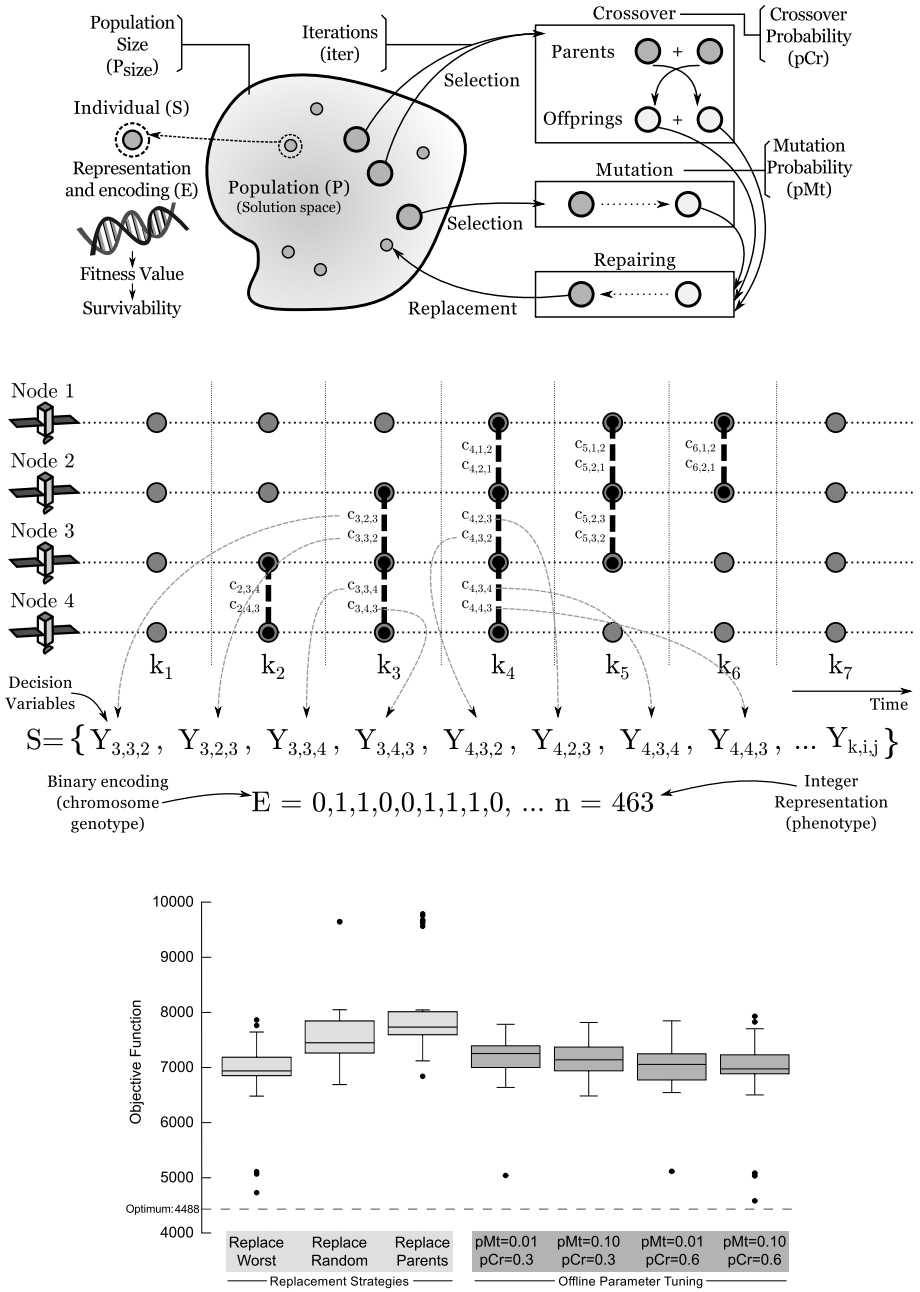

Figure B.13: Traffic-Aware Contact Plan based on Evolutionary Algorithms (TACP-EA). **Top Panel**: Illustration of the core functions used in the GA approach for traffic-aware contact plan design. The process includes selection (choosing parent solutions based on fitness), crossover (combining parent solutions to create offspring), mutation (introducing random changes for diversity), and repairing (ensuring offspring feasibility by addressing any constraint violations). **Center Panel**: Graphical representation of the GA encoding technique specifically developed for the contact plan design problem. Each individual (solution) is encoded as a binary chromosome (*genotype*) representing decision variables ($Y_{k,i,j}$), which are subsequently interpreted into an integer representation (*phenotype*). **Bottom Panel**: Results of hyperparameter tuning for TACP-EA, comparing the impact of different replacement strategies (e.g., replacing the worst, random, or parent solutions) and parameter tuning for mutation (*pMt*) and crossover (*pCr*) probabilities. (From [Fra+17c])

Meeting and collaborating with Prof. Holger Hermanns after completing my PhD was a transformative milestone in my career. His guidance and his team's work introduced me to rigorous computer science methodologies, significantly enhancing and complementing my engineering expertise in ways I continue to uncover. I am deeply grateful to Holger, now a close friend and invaluable mentor, whose wisdom and countless anecdotes enrich every interaction.

BACP In collaboration with Prof. Hermanns' team and GomSpace, a Danish company with operations in Luxembourg, we explored the intricacies of designing contact plans for distributed nano-satellite missions with constrained battery capacities, or Battery-Aware Contact Plan Design (BACP). Preliminary findings published in [Fra+18b] and further expanded in the journal article [Fra+20b] tackled the pressing challenge of crafting battery-aware contact plans for LEO satellite constellations. Often operating under stringent power constraints, these networks risk premature battery depletion, particularly during prolonged eclipse periods. Our research on BACP focused on incorporating battery-aware scheduling into the contact plan design process, addressing the need to balance efficient communication with sustainable energy utilization and ensuring the longevity and reliability of the satellite network.

For BACP, we proposed a MILP model incorporating communication and battery dynamics. The MILP formulation includes constraints on battery charge levels, transmission power consumption, and energy replenishment during sunlight exposure. Two battery models are considered: the simplistic Linear Battery Model (LiBaM) and the more accurate but computationally intensive Kinetic Battery Model (KiBaM). The MILP model uses LiBaM for optimization and validates its output with KiBaM in a post-processing step to ensure the robustness of the contact plan. The proposed framework is applied to the Ulloriaq case study, a satellite constellation inspired by the GOMX-4 Mission, consisting of 10 satellites in a train-like configuration with inter-satellite and satellite-to-ground communication links. A graphical depiction of the Ulloriaq constellation and the LiBaM and KiBaM models is presented in Figure B.14.

The battery-aware contact plans generated by the MILP model significantly reduce the risk of battery depletion compared to battery-agnostic plans. As illustrated in Figure B.14, while battery-agnostic schedules result in certain satellites depleting their batteries entirely, the battery-aware approach reduces depletion risk to negligible levels. The study demonstrates that the battery-aware method slightly reduces overall data throughput compared to battery-agnostic schedules but ensures network sustainability by evenly distributing communication activities across time and nodes.

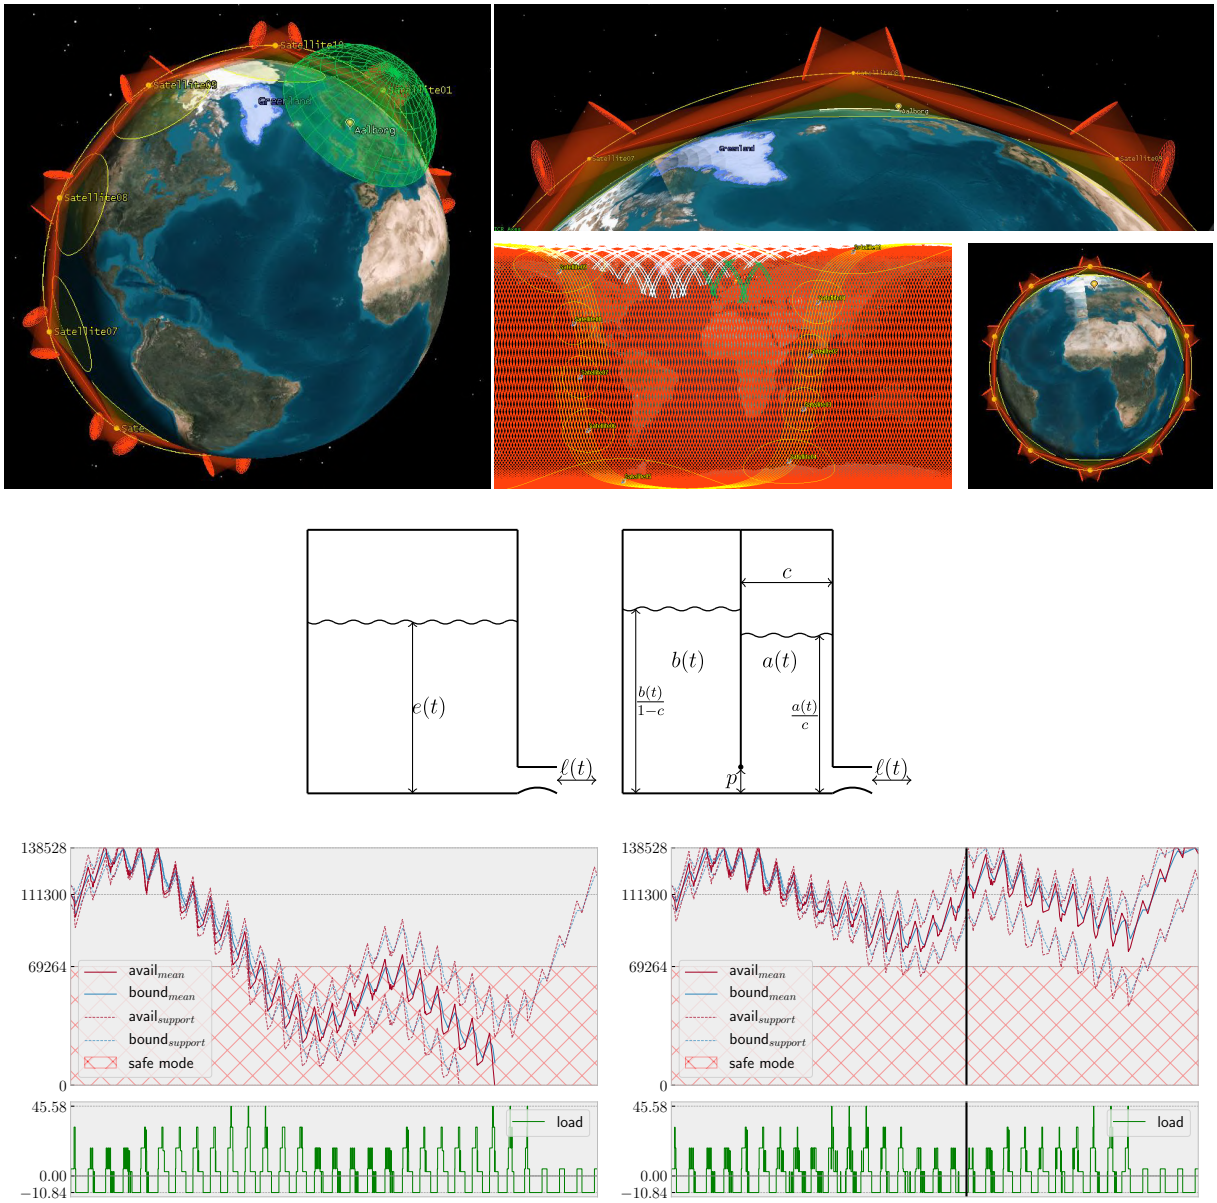

Figure B.14: Battery-Aware Contact Plan Design applied to the Ulloriaq constellation. **Top Panel**: A STK model of the Ulloriaq constellation comprising 10 satellites, equally spaced in the same orbital plane, with Aalborg designated as the primary ground station. Central panel: The LiBaM and KiBaM battery models illustrate the energy dynamics of satellite operations. The Linear Battery Model assumes a simplified linear discharge behavior. In contrast, the Kinetic Battery Model incorporates dual energy reservoirs (available $a(t)$ and bound $b(t)$) to represent the interplay between immediate and constrained energy availability. **Bottom Panel**: The state-of-charge evolution for a satellite is shown under a battery-agnostic design (left) and a battery-aware design (right). The battery-agnostic plan leads to premature depletion during high-load phases, demonstrating the risks of ignoring battery dynamics. In contrast, the battery-aware plan distributes the load evenly across contacts, reducing depletion risks. Although the battery-aware worst-case scenario still approaches critical depletion levels, the load balancing and battery health improvements are significant compared to the agnostic plan. (From [Fra+20b])

Scalability Building on the battery-aware CPD discussed in previous work, which relied on MILP models but suffered from scalability issues, my paper in [Fra+21] extends the framework to address the operational challenges posed by prominent constellations. For this work, I received support from Carsten Gerstacker, a Master's student I supervised at the Saarland University. To improve scalability, we introduced enhancements to the MILP formulation, such as refining battery constraints with a quantitative representation of link usage, deriving a pure Linnear Programming (LP) for no interface constraints, and relaxing assumptions that previously over-approximated energy consumption. These enhancements allowed a 60% increase in the size of solvable problem instances compared to earlier models. The study extensively evaluated benchmarked constellations, including single and dual-ring configurations with up to 50 satellites. As detailed in Figure B.15, results reveal substantial enhancements in computation times and memory efficiency for scenarios with limited scheduling horizons. These optimizations enabled the proposed models to deliver tractable solutions for constellations comprising up to 40 satellites, showcasing their scalability and applicability to larger, more complex satellite networks. However, the paper also highlights that current techniques remain insufficient for the day-to-day operation of mega-constellations with hundreds or thousands of satellites. It emphasizes further research into heuristic-based or hybrid approaches to extend scalability.

B.2.2.5 Position-Awareness

This section highlights my contributions to a novel optimization criterion for the CPD problem, focusing on the relative positioning of space nodes when selecting contacts. Relative positioning plays a critical role in Global Navigation Satellite System (GNSS) systems, where achieving positioning diversity is essential to enhance the accuracy of positioning determination. In this context, I collaborated with Zhibo Yan, a PhD student I co-supervised alongside Prof. Kanglian Zhao from Nanjing University, China. Additionally, another project stemmed from a fruitful collaboration with Andrea Nardin, a PhD student under the supervision of Prof. Fabio Dovis at the Politecnico di Torino, Italy. These works, published in IEEE Transactions on Aerospace and Electronic Systems journal, focus on two prominent GNSS systems: **BeiDou GNSS** (China) and **Kepler GNSS** (European Union), addressing system-specific challenges and advancing the state of positioning-aware contact plan design.

BeiDou GNSS The paper in [Yan+20] addresses the problem of designing a distributed contact plan for next-generation GNSS equipped with ISLs. Traditionally, GNSS systems rely on centralized control for CPD, leading to high dependency on ground stations, increased operational costs, and reduced outage robustness. To address this, our paper

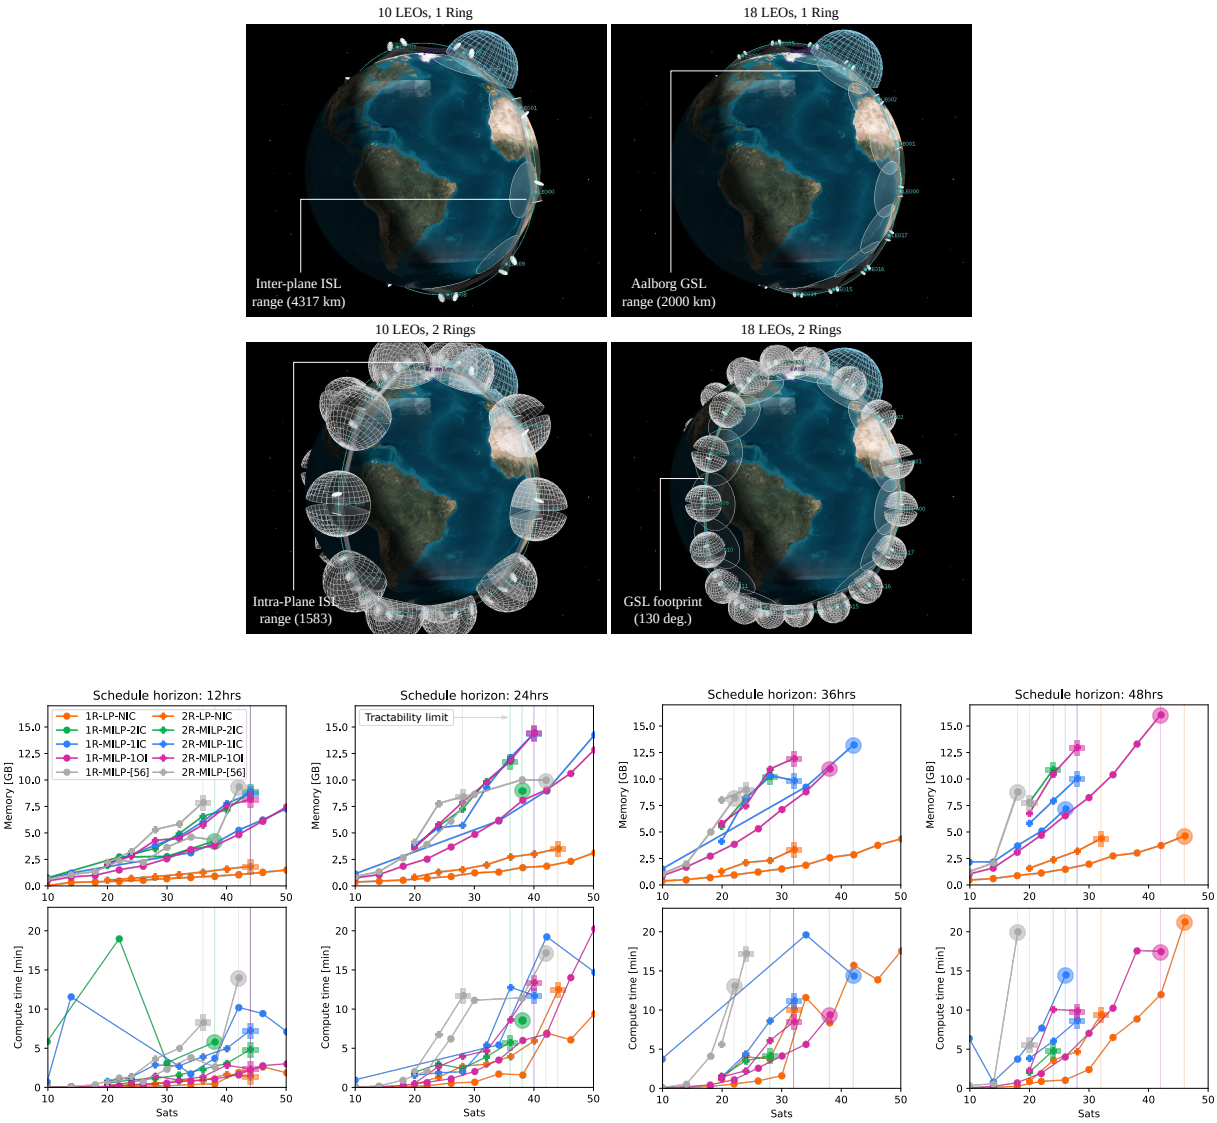

Figure B.15: Battery-aware CPD scalability analysis. **Top Panel**: LEO satellite constellations evaluated for battery-aware contact plan design. The top-left panel shows a single-ring configuration with 10 satellites interconnected by inter-plane ISL with a maximum range of 4317 km. The top-right panel depicts an 18-satellite single-ring configuration. The bottom-left panel features a two-ring configuration with 10 satellites per ring, where intra-plane ISLs have a shorter range of 1583 km. The bottom-right panel illustrates a two-ring configuration with 18 satellites per ring, emphasizing the extended ground station coverage footprint with a 130° field of view. **Bottom Panel**: Metrics obtained from solving various optimization models: LP-NIC (Linear Programming with No Interface Constraints), MILP-1IC (Mixed-Integer Linear Programming with 1 Interface Constraint), MILP-2IC (2 Interface Constraints), MILP-1OI (1 Outgoing Interface Constraint, representing the new model), and MILP-[56] (1 Outgoing Interface Constraint, representing an older model). These models are evaluated for the 1-ring (1R) and 2-ring (2R) benchmark scenarios. Tractability limits are indicated with larger markers and vertical lines, highlighting the latest point where each model rendered a solution within 20 minutes. Notably, configurations for 1-ring scenarios begin with 10 satellites, while 2-ring scenarios involve a minimum of 20 satellites (10 per ring).(From [Fra+21])

introduced the first distributed CPD scheme, the Distributed Fair Contact Plan (DFCP), which allows satellites to autonomously compute their contact plans without relying on centralized infrastructure. DFCP is built upon a deterministic algorithm derived from the Fair Contact Plan Design (FCP) framework, adapted to optimize both ranging via Position Dilution of Precision (PDOP) and data delivery delay metrics. The proposed method uses an efficient non-perfect weighted matching algorithm to dynamically assign Inter-Satellite Links (ISL)s, enabling decentralized coordination while preserving global consistency. Key outcomes demonstrate that DFCP achieves performance metrics comparable to state-of-the-art centralized schemes, with significant autonomy and computational efficiency advantages. DFCP provides accurate PDOP and competitive delay metrics while requiring only a fraction of the computational resources, validating its feasibility for deployment on resource-constrained onboard satellite hardware. This work pioneer distributed CPD for GNSS, marking a significant step toward autonomous satellite network operations.

Kepler GNSS Our paper in [NFD22] addresses the problem of optimizing contact plans for the Kepler GNSS constellation, which leverages Optical Inter-Satellite Link (OISL)s to enhance orbit determination and clock synchronization while reducing ground station dependency. The proposed method integrates Degree-Constrained Minimum Spanning Tree (DCMST) algorithms to design a contact plan that minimizes the PDOP, a critical metric for positioning accuracy. The study considers the constraints imposed by the limited number of steerable laser communication transceivers and their Field of Regard (FoR), as well as the dynamic connectivity requirements among Medium-Earth Orbit (MEO) satellites and a supporting LEO backbone. The method employs a graph-based approach, where PDOP values are edge weights. The DCMST ensures network connectivity while respecting transceiver constraints. The algorithm is further extended with secondary optimization strategies to maximize transceiver utilization and prioritize metrics such as revisit time (RT) or multi-plane connections. Key outcomes demonstrate significant improvements in PDOP reduction, with the optimal algorithm achieving up to 85% better positioning accuracy compared to non-optimized schedulers. Additionally, the proposed algorithms maintain fairness in resource allocation and achieve tractable computational times for constellation scales typical of GNSS systems, paving the way for scalable deployment in real-world Optical GNSS networks. The Kepler case study validates the approach, showcasing how tailored contact planning significantly enhances navigation and data transfer performance.

LEO-Based GNSS Also developed with Andrea Nardin, the paper in [NDF20] addresses the challenge of leveraging LEO satellite constellations for Positioning, Navigation, and

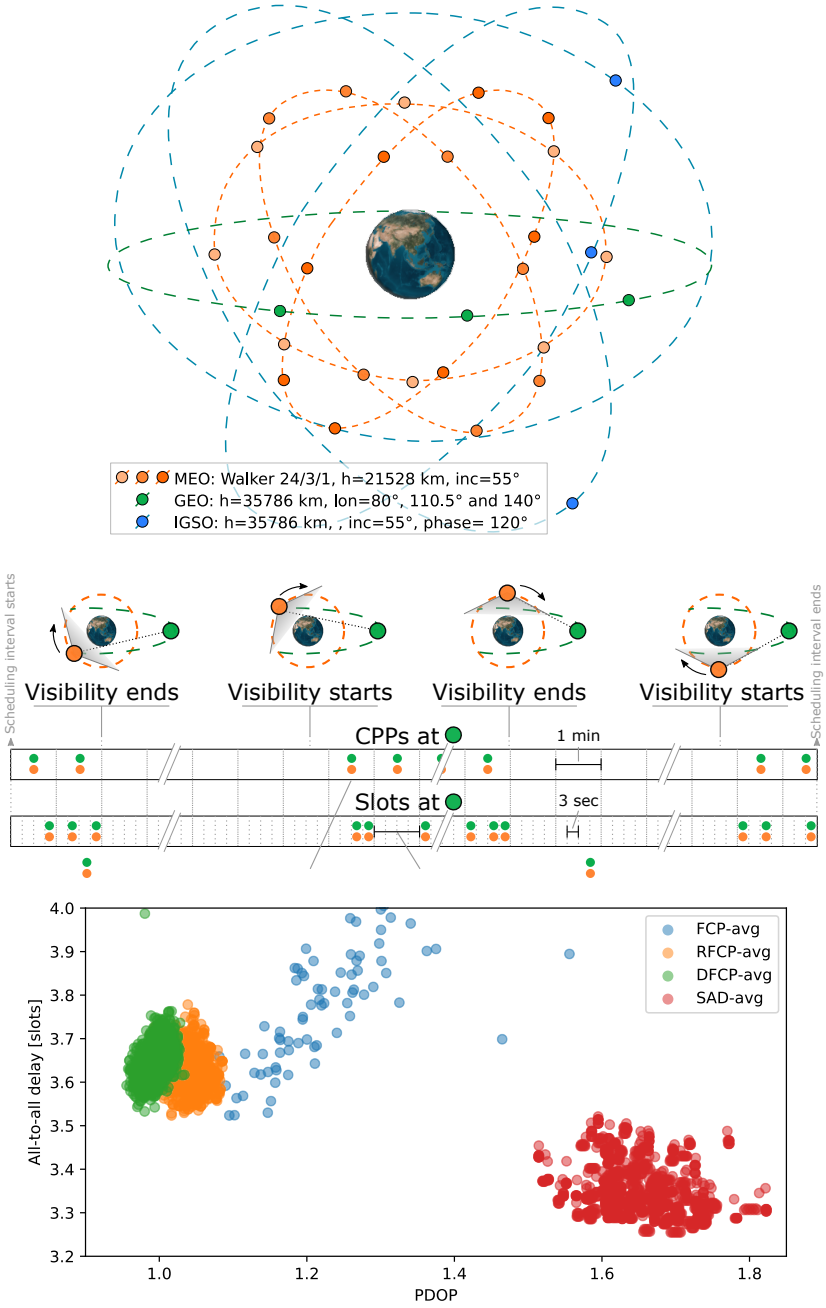

Figure B.16: Distributed Fair Contact Plan (DFCP) in BeiDou GNSS. **Top Panel**: The orbital configuration of the BeiDou GNSS constellation includes three types of satellites: Medium-Earth Orbit (MEO) satellites, GEO, and Inclined Geosynchronous Orbit (IGSO) satellites. **Center Panel**: The GNSS topology model is represented through Contact Plan Periods (CPP) and slots. Contact Plan Periodss define fixed intervals during which inter-satellite and satellite-to-ground visibility are available. At the same time, slots divide these periods into finer time granularity for scheduling communication and positioning tasks. **Bottom Panel**: A scatter plot compares average all-to-all delay (in slots) and PDOP metrics for four CPD schemes: FCP, Route-Aware Contact Plan Design (RACP), DFCP, and Simulated Annealing Distributed (SAD) plan. DFCP achieves a balance between low PDOP and reduced delay, whereas SAD offers the best performance for both metrics. (From [Yan+20])

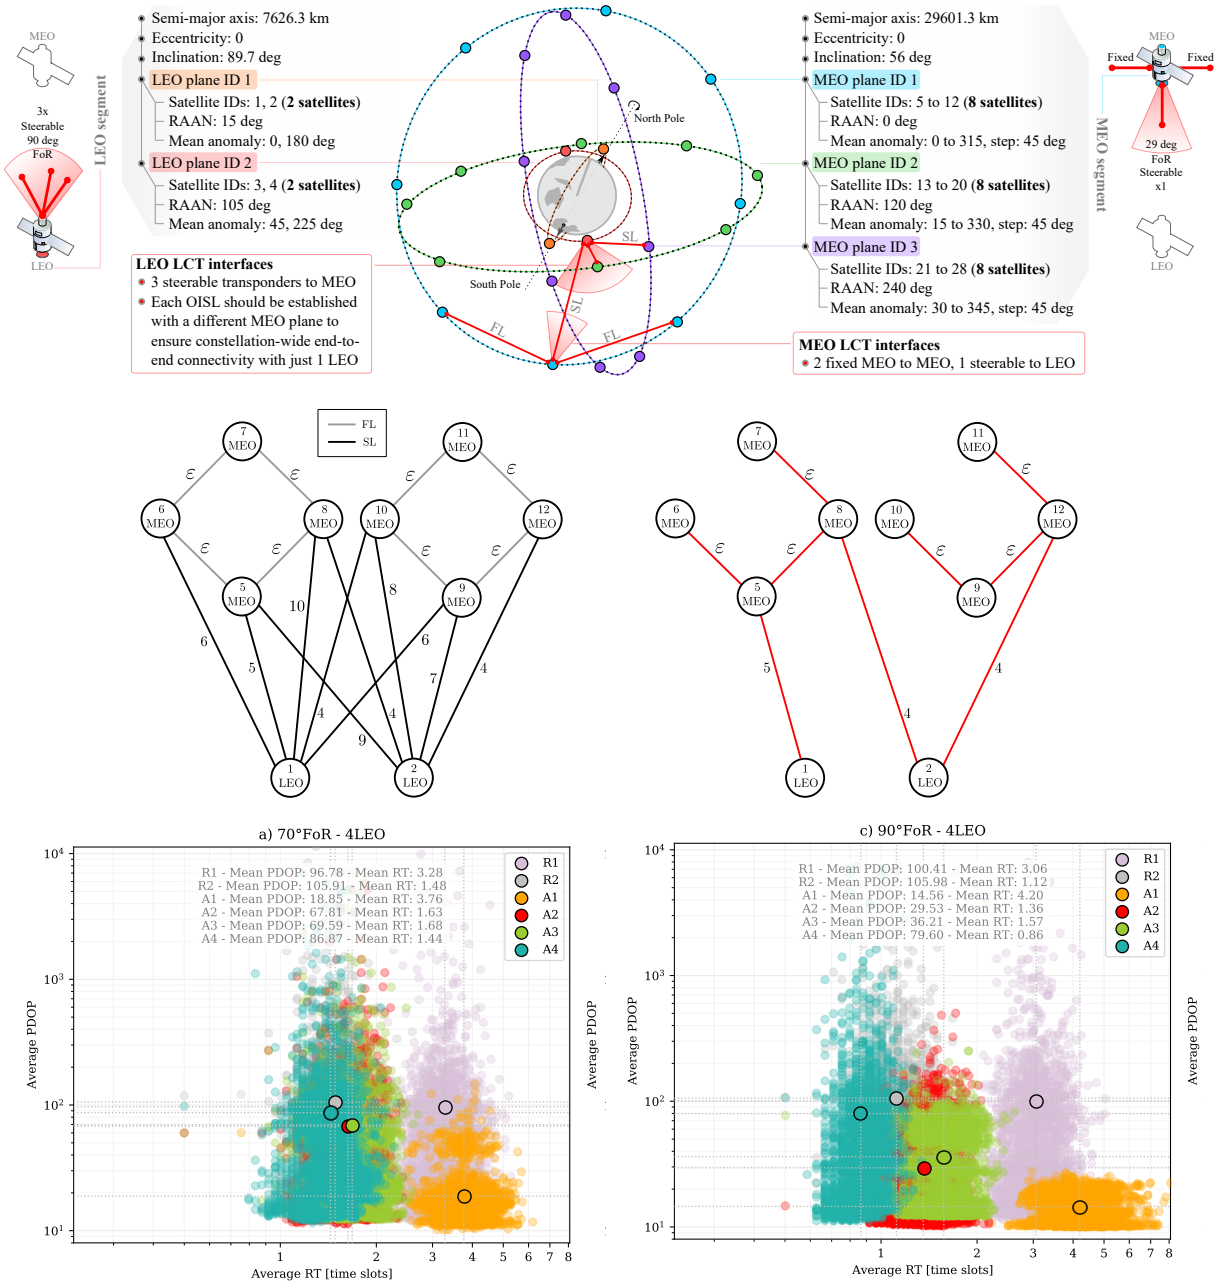

Figure B.17: Degree-Constrained Minimum Spanning Tree (DCMST) in Kepler GNSS. **Top Panel**: The Kepler GNSS system architecture, featuring two orbital layers: LEO and Medium MEO. ISLs are employed for communication within the LEO layer, as well as between the LEO and MEO layers. **Center Panel**: Visualization of the DCMST algorithm in practice. The algorithm selectively retains a subset of edges in the connectivity graph based on PDOP as the core metric. FL refers to the Fixed Links between MEO satellites, while SL represents Steerable Links connecting LEO to MEO. **Bottom Panel**: Scatter plots depicting the results of PDOP and average revisit time (RT) for various methods. The methodologies include A1, A2, A3, and A4 (four DCMST variants) and R1 and R2 (two DCMST termination conditions), described in detail in the paper. (From [NFD22])

Timing (PNT) services while overcoming limitations such as high Doppler effects and constrained bandwidth, which are inherent to these systems. Although this is not strictly a CPD contribution, this work helped me to understand the relevance of highly dynamic positioning and PDOP computations, which are much more complicated than in MEO. The study investigates the implementation of a Virtual Wideband (VWB) meta-signal processing strategy. This strategy combines multiple narrowband PNT signals to emulate a wideband signal, thus improving code tracking accuracy and reducing estimation errors under high dynamics. We simulated a fused PNT service on Amazon's Kuiper LEO constellation and generated realistic BPSK-modulated signals subjected to Doppler profiles and atmospheric conditions. A custom GNSS software receiver employing VWB processing is then used to analyze the generated signals. Key outcomes demonstrate that VWB processing significantly reduces root mean square errors (RMSEs) of code-rate estimations, particularly under high Doppler conditions. In challenging scenarios, VWB processing achieved an RMSE reduction of over 70%. The findings validate the potential of VWB architectures to empower high-accuracy PNT services using LEO satellites while maintaining efficiency under dynamic conditions.

B.2.2.6 Tooling: Contact Plan Designer

I have had the opportunity to engage with various non-academic organizations and personnel. For instance, I collaborated with GomSpace on the Ulloriaq and GomX-4 nanosatellite missions, which utilized cross-linked architectures, and with Skyloom, focusing on optical inter-satellite links (ISL) between LEO and GEO satellites, where pointing decisions directly influenced contact plan selections. My early work with CONAE, the Argentinian Space Agency, on the SARE mission also addressed the planning complexities of limited ISLs. I recognized this gap and developed a tool combining strong visual representation with technical precision to effectively communicate and interact with CPD. The result was the aptly named "Contact Plan Designer."

> In my interactions, explaining Contact Plan Design (CPD)—what it is, how it is implemented in real-world space networks, why it is critical, and its impact on overall system performance—proved exceptionally challenging, especially when communicating with non-technical stakeholders, such as decision-makers.

The tool was presented in the paper in [Fra17]. It addresses the challenge of creating precise, resource-optimized contact plans for space DTNs. The CPD tool automates contact plan design, analysis, and optimization, accounting for constraints such as limited transceivers, power budgets, and interference. The proposed tool integrates with Systems Tool Kit (STK), a validated physics-based software, to accurately model

and predict communication opportunities between space and ground assets. CPD's workflow includes manual and automated methods for resolving conflicts in contact topologies, optimizing the contact schedule, and enabling comprehensive network performance analysis, including routing and queuing behavior. A proof-of-concept scenario involving two LEO satellites and a ground station demonstrated CPD's ability to generate actionable contact plans for realistic constraints like line-of-sight and link budgets. The outcomes highlight the CPD Tool's potential as a foundational tool for operational space network deployments, bridging the gap between academic advancements and real-world implementation. Future iterations incorporated advanced routing and traffic-aware mechanisms, further solidifying this tool's role in next-generation space communication systems.

B.3 Mission Operation

The field of satellite mission operations is undergoing a transformative shift. Advanced automation and optimization techniques are replacing traditional manual interventions. This evolution is driven by the increasing complexity of modern constellations and the demand for energy-efficient, scalable, and autonomous operations. My research in mission operations bridges the domains of computer science and space engineering, focusing on developing practical tools and methodologies that address real-world challenges. From validating automated, battery-aware operations in the GomX-4 mission to advancing mission planning through precise coverage models, network size estimation techniques, and satellite visibility prediction algorithms, these contributions have exemplified the integration of theoretical advancements into practical systems.

B.3.1 Automated Battery-Aware Operations

One of the defining moments in my research journey was the opportunity to validate advanced mission operations concepts in orbit, transitioning theoretical work into tangible outcomes. This effort was exemplified by the collaboration with GomSpace, which graciously allowed access to their dual-satellite mission, GomX-4, for experimental purposes. The mission, comprising two 6U CubeSats equipped with ISLs, was a proving ground for validating battery-aware Contact Plan Design (CPD) developed during my PhD (discussed in Section B.2.2.4) and extended thereafter for this experiment with GomX-4 A and GomX-4 B nano-satellites. These designs sought to automate the operation of satellite networks while ensuring energy-efficient utilization of constrained resources, such as onboard batteries. The core of this effort was carried out by Gregory Stock, a PhD student I co-supervised with Prof. Holger Hermanns.

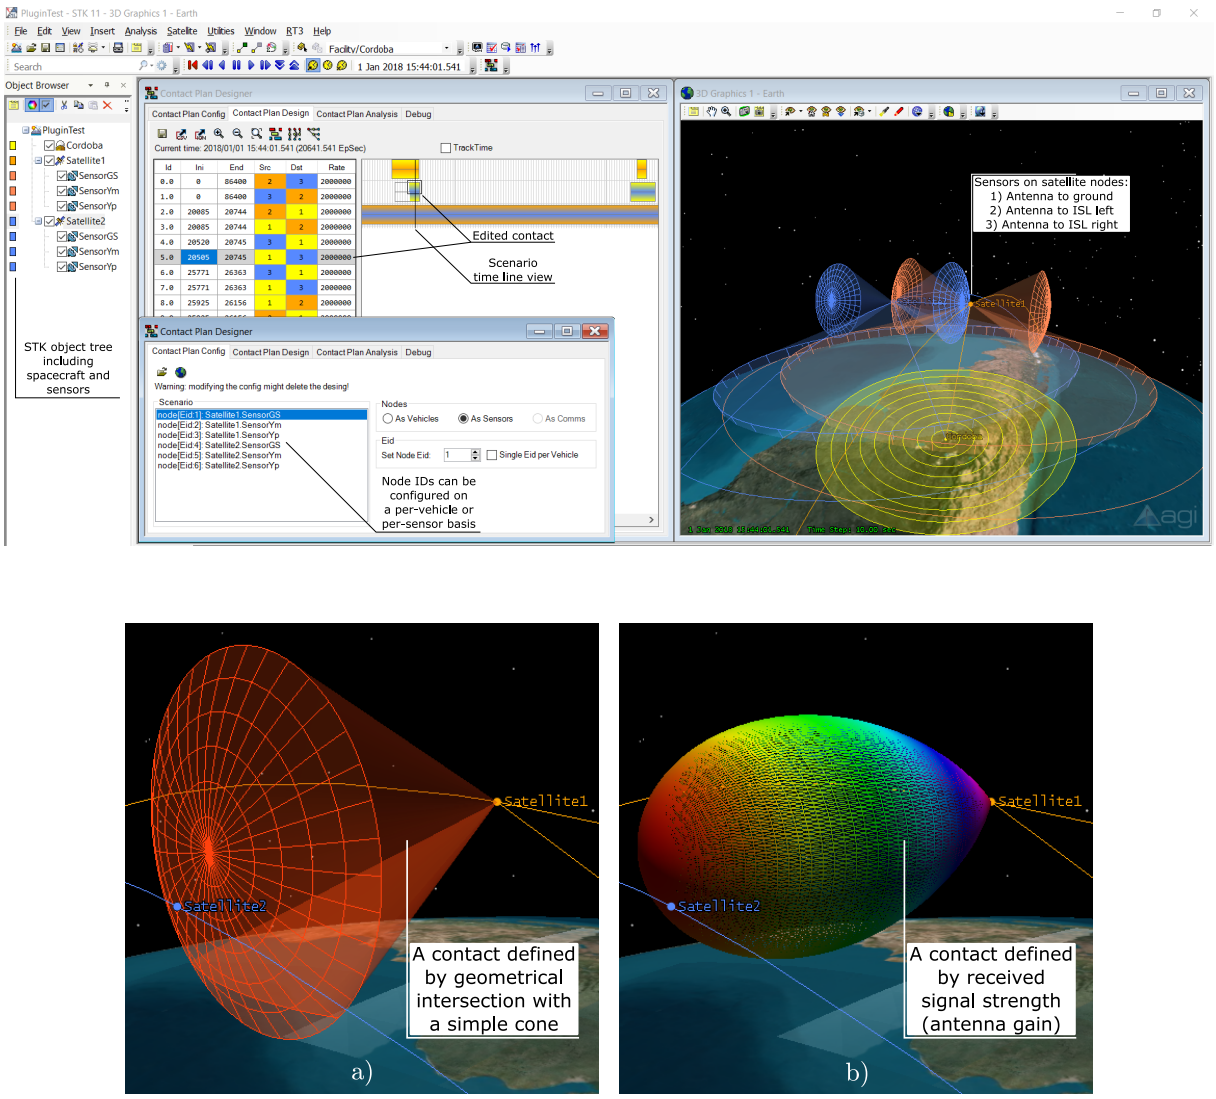

Figure B.18: Core Features of the Contact Plan Designer Tool. **Top Panel**: The Contact Plan Designer's graphical user interface integrates with the STK environment. The left window displays the object tree, which organizes spacecraft and sensor nodes. The tool also offers flexibility regarding the node identifiers, which can be set to spacecraft, transponders, or antennas. The central timeline view provides a scenario editor to modify and visualize contact schedules, highlighting the interaction between contacts and timeline configurations. The 3D view on the right visualizes satellite nodes and sensors, illustrating their antenna roles: ground communication, left ISL, and right ISL connections. **Bottom Panel**: Demonstration of the tool's flexibility in antenna modeling. (a) A contact is defined using simple geometric cones to establish a link based on intersection criteria. (b) A contact is modeled using signal strength contours derived from antenna gain, offering enhanced realism and precision in contact definitions. This flexibility allows users to tailor the CPD configuration to diverse mission requirements and capabilities. (From [Fra17])

We collaborated with Saarland University colleagues and developed and integrated advanced software tools that leveraged Dynamic programming (DP), antichain-based pruning, and non-linear battery models. The tool was presented at IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems journal [Sto+20], co-authored with GomSpace staff that supported our experiments. These tools ensured the optimal and automated scheduling of payload and communication tasks while maintaining robust operation under real-world constraints. The platform managed solar energy harvesting and payload operations and incorporated *receding-horizon scheduling* to adapt plans dynamically based on real-time telemetry (see Figure B.19). The results demonstrated the feasibility of operating a constellation without human intervention while achieving energy efficiency and optimal task scheduling.

Key outcomes of this work included the development of the LEOPOWVER toolchain (see Figure B.19) and the project website for further details), which interfaced seamlessly with GomSpace's Hands-Off Operations Platform (HOOP). By integrating Kinetic Battery Model (KiBaM) and satellite telemetry data, the system demonstrated its ability to optimize task execution while mitigating risks of battery depletion. A notable highlight was the fully automated validation campaign conducted over a 48-hour in orbit, where the system operated autonomously to maximize mission objectives. These experiments validated the scalability and robustness of our methodologies and set a precedent for autonomous satellite operation in modern constellations. The results were very well taken by the industrial space community, as highlighted in our invited paper to the highly popular SmallSat conference [Sto+22].

This body of work showcased the potential of applying advanced computer science techniques to mission operations. It moved beyond manual intervention to a model where operators focus on strategic objectives. Integrating formal methods, efficient scheduling algorithms, and real-world validation created a feedback loop between theoretical advancements and practical applicability. The experience gained from the GomX-4 mission laid a foundation for future collaborations and established a benchmark for automated satellite constellation management in academic and industrial domains.

The successful **demonstration of automated satellite operations** garnered significant media attention, placing myself, Prof. Holger Hermanns ' team, and GomSpace in the spotlight within the space community for several weeks (e.g., see SIC News). This fruitful collaboration with GomSpace later led to their involvement in the MISSION EU project, where we continue to build upon and refine this groundbreaking experience.

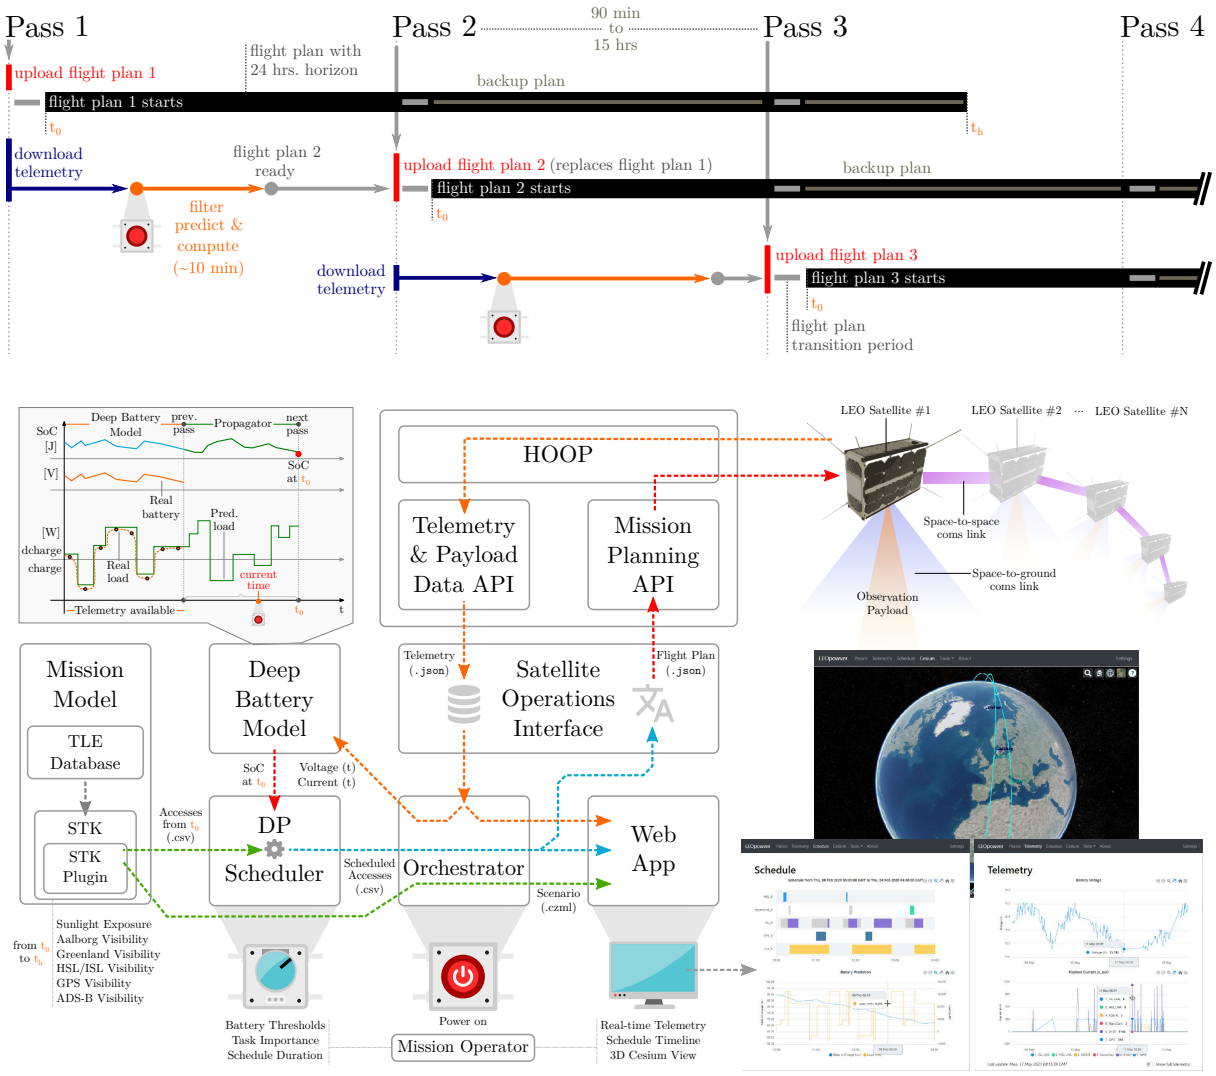

Figure B.19: GOMX-4 Mission: Automated Battery-Aware Operations. **Top Panel**: The receding horizon approach dynamically adjusts flight plans based on real-time telemetry and predicted battery states of charge (SoC). Each satellite pass involves uploading a flight plan with a predefined time horizon, downloading telemetry from the previous plan, and processing this data to refine subsequent plans. This iterative process ensures continuous optimization of battery and mission operations while incorporating updated environmental and system conditions. **Bottom Panel**: The LEOPOWVER toolchain integrates mission modeling, deep battery models (e.g., KiBaM), Dynamic programming (DP) schedulers, and an orchestrator connected to the GomSpace Hands-Off Operations Platform (HOOP). The system supports automated satellite operations with telemetry, payload, and mission planning APIs. It includes a web-based interface for real-time telemetry visualization, schedule adjustments, and 3D constellation views, enabling efficient and energy-aware task execution in cross-linked satellite networks. (From [Sto+20] - More information in LEOPOWVER website)

B.3.2 Advanced Operation Models

The experience with the GomX-4 mission was a pivotal milestone in my research, solidifying that mission operations in modern satellite constellations are undergoing a transformative evolution. Leveraging advanced automation and optimization techniques, these operations increasingly reduce human intervention while maximizing resource efficiency. Building on this foundation, my work has continued to bridge computer science and space engineering, delivering innovative solutions for satellite operations. This section highlights additional key achievements, including developing and validating mission operation frameworks with a focus on **polygon-based coverage models**, **network size estimators**, and **satellite visibility pedictors**. These contributions exemplify the tangible impact of theoretical advancements and underscore the potential of interdisciplinary collaborations that drive my research forward.

B.3.2.1 Polygon-Based Coverage Models

Led by Santiago Henn, my PhD student at CONICET, our paper in [HFH23] addresses the complex problem of calculating satellite coverage metrics for LEO constellations over arbitrary Regions of Interest (ROI) on Earth, traditionally handled by net-point division techniques. Existing methods' limitations, such as their reliance on spherical Earth approximations, fixed grids, and approximate results, are overcome by leveraging polygon-based Boolean operations and transformations between Euclidean and non-Euclidean spaces. The proposed method models satellite access regions as precise, non-self-intersecting polygons. It leverages fast orbital propagators and transformations to project these regions into a 2D Euclidean plane for efficient intersection analysis. This approach avoids grid-based approximations and integrates Earth's oblateness to model satellite footprints more accurately. The results demonstrate significantly improved accuracy and computational efficiency compared to state-of-the-art tools like STK, especially in scenarios requiring fine-granularity coverage analysis. The methodology supports diverse mission requirements, including Earth observation, communication, and navigation, while enabling precise assessment of coverage dynamics over time.

B.3.2.2 Network Size Estimators

Optimistic Collision Information Our paper in [Par+23], led by Pablo Illabaca, a Master student I hosted at Inria and co-supervised with Prof. Sandra Céspedes, addresses the critical problem of network size estimation in DtS-IoT networks, where thousands of constrained IoT devices must communicate through LEO CubeSat gateways with

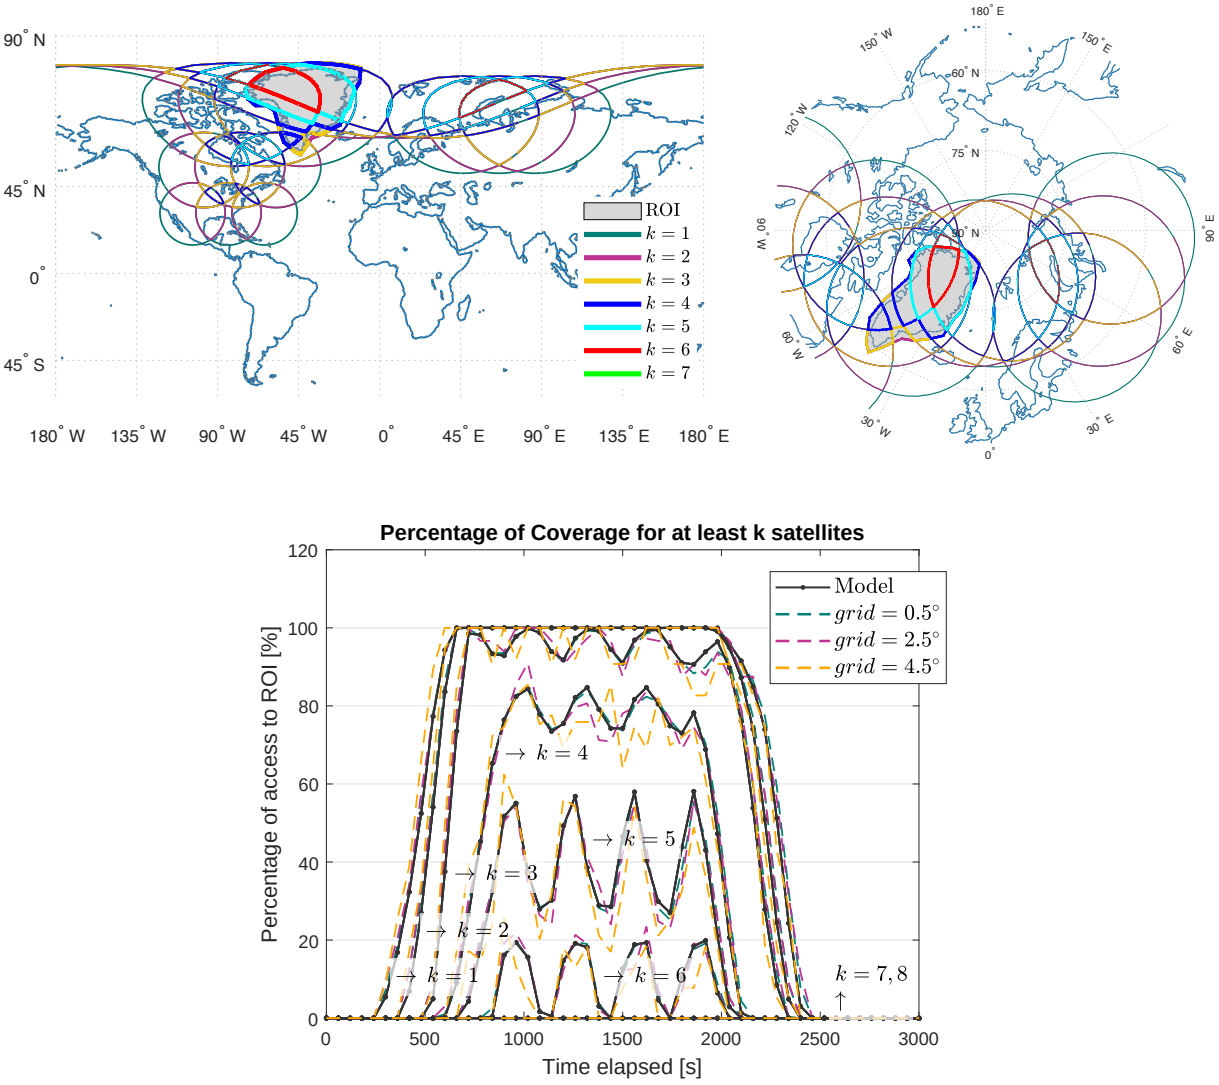

Figure B.20: Polygon-Based Coverage Analysis Algorithms. **Top Panel**: Visualization of polygon transformations representing satellite coverage regions for a given ROI. The color-coded regions indicate the number of overlapping satellite coverages (k), where k=1, 2, . . . , 7. The singularity at the poles is accurately managed through non-Euclidean-to-Euclidean transformations, ensuring precise modeling even in high-latitude regions. **Bottom Panel**: Percentage of coverage over time for each k value, comparing the proposed polygon-based model ("Model") against grid-based methods at varying resolutions (0.5°, 2.5°, and 4.5°). The results demonstrate the high accuracy of the polygon-based approach, which closely matches the performance of fine-resolution grids but achieves significantly faster computation. For this scenario, the polygon-based model delivered results in 13.2 seconds, compared to the 49.7 seconds required by the finest grid-based method in STK, highlighting its computational efficiency. (From [HFH23])

limited computational and energy resources. The main challenge lies in accurately estimating the number of active IoT nodes in the satellite's footprint to optimize MAC protocols and enhance system scalability. To tackle this, we proposed a novel, computationally efficient network size estimation mechanism termed the Optimistic Collision Information (OCI)-based estimator. OCI uses a lightweight polynomial correction applied to observed collision data from Frame-Slotted Aloha (FSA) communication windows to estimate the node population. The method avoids the computational complexity of traditional maximum likelihood estimators by leveraging simplified calculations suitable for constrained CubeSat hardware. Extensive simulations demonstrate that leveraging OCI in a Slotted ALOHA Game (SAG) achieves high scalability, supporting networks with up to 4,000 devices while maintaining low estimation error and energy efficiency. Compared to state-of-the-art alternatives, OCI reduces computation time by up to 75%, making it a viable candidate for real-world DtS-IoT applications. Results of using OCI's feedback in a Slotted ALOHA Game (SAG) access are summarized in Figure B.21. The paper further shows that OCI-driven MAC protocols achieve throughput near theoretical limits, offering an 8-fold improvement in scalability over pure FSA schemes.

OCI for LoRa/LoRaWAN The paper in [Mal+23], led by Diego Maldonado, my PhD student builds on the OCI estimator framework to adapt it for the specifics of LoRa/LoRaWAN networks in DtS-IoT systems. The primary problem addressed involves the scalability and accuracy of network size estimation for thousands of ground IoT nodes communicating directly with a LEO satellite. The OCI model in [Par+23], while efficient, lacked compatibility with the unique synchronization and channel characteristics of LoRa networks. The authors propose LoRa-OCI (L-OCI), which extends OCI to accommodate LoRa's specific modulation schemes and time-slotted communication protocols (e.g., LoRaSync). L-OCI incorporates a multi-objective optimization algorithm to balance energy consumption and estimation accuracy by fine-tuning LoRa parameters such as Spreading Factor (SF), Bandwidth (BW), and beacon length. Implemented and evaluated using the FLoRaSat simulation platform, L-OCI demonstrated high accuracy in node estimation with low Root Mean Square Error (RMSE) across diverse scenarios, achieving throughput optimization in Slotted ALOHA Game (SAG)-based MAC protocols. The parameter tuning campaign, summarized in Figure B.21, identified configurations (e.g., SF9, BW125kHz) that balanced power efficiency and RMSE, suggesting minimal performance trade-offs when resources are constrained. L-OCI presents a practical and energy-efficient solution for integrating network size estimation into LoRa-based DtS-IoT frameworks, enabling enhanced MAC protocol performance in large-scale deployments.

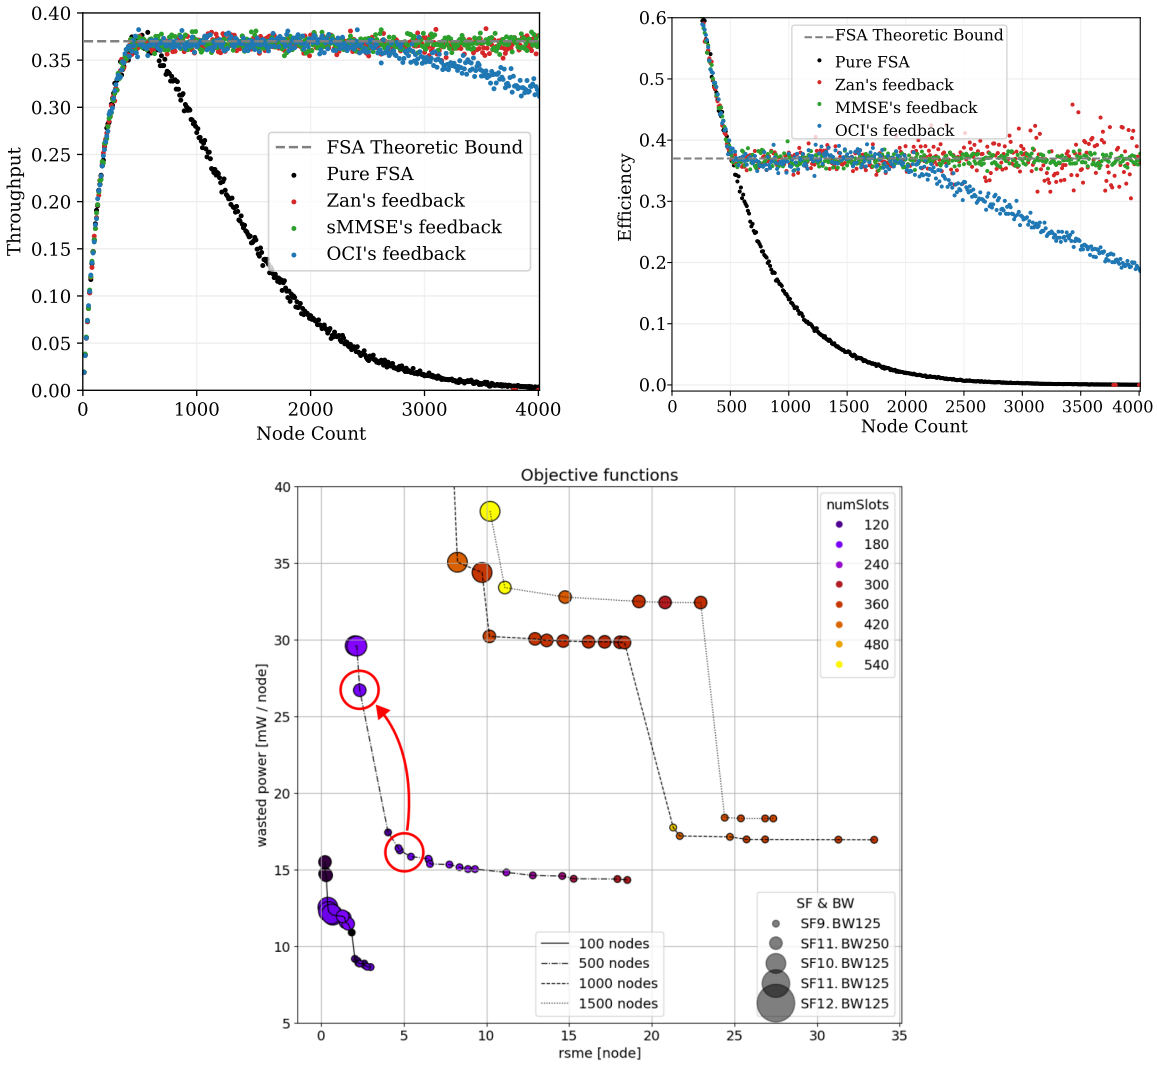

Figure B.21: Network Size Estimation: using OCI's feedback for an Slotted ALOHA Game (SAG). **Top Left Panel**: Throughput performance as a function of the node count. The dashed line represents the FSA Theoretic Bound. The Pure FSA curve (black dots) shows a significant drop in throughput as the node count increases due to collisions. Zan's feedback (red dots) and MMSE's (green dots) demonstrate throughput improvements by integrating network size estimators to optimize retransmission probabilities. In contrast, the proposed OCI's feedback (blue dots) closely approaches the theoretical bound, maintaining stable performance even at high node densities. **Top Right Panel**: Efficiency as a function of the node count. Efficiency reflects the fraction of successfully transmitted slots versus the total slots used. The Pure FSA curve drops sharply as node count increases, indicating poor utilization of the communication window under high contention. Zan's and MMSE's feedback shows efficiency improvements but deviates under heavy loads. OCI's feedback, leveraging its computational simplicity and accuracy, maintains higher efficiency across all node densities. (From [Par+23]) **Bottom Panel**: Network Size Estimation: OCI adaptation for LoRa/LoRaWAN. The trade-off between wasted power per node (mW/node) and Root Mean Square Error (RMSE) in network size estimation under various configurations of spreading factor (SF), bandwidth (BW), and the number of time slots (numSlots). (From [Mal+23])

OCI for LoRa/LoRaWAN Constellations Our recent paper in [Ila+25] introduces Constellation LoRa Optimistic Collision Information (CL-OCI), a method for estimating the network size in scenarios involving multiple LEO satellites. The method divides the network into groups to reduce contention and energy consumption, incorporates adjustments for satellite ground track shifts due to Earth's rotation, and aggregates individual satellite estimates into a constellation-wide result. Simulations using the FLoRaSat platform demonstrate that CL-OCI doubles the maximum number of nodes accurately estimated compared to single-satellite approaches while reducing energy wasted due to collisions by seven times. This work showcases a scalable and energy-efficient solution for resource allocation and Medium Access Control (MAC) in large-scale DtS-IoT networks.

B.3.2.3 Satellite Visibility Predictors

Our paper in [Ort+24], led by Raydel Ortigueira, a PhD student I co-supervised with Prof. Sandra Céspedes, addresses the challenge of accurate and efficient satellite visibility prediction for constrained DtS-IoT devices. These devices must conserve energy while maximizing data transmission opportunities during short visibility intervals of LEO satellites. To tackle this problem, we proposed a lightweight algorithm based on the two-body orbital model that estimates satellite visibility intervals while incorporating an innovative mechanism for ground terminals to update orbital parameters without internet connectivity. The method includes iterative refinements for precise visibility determination and leverages sync packets to synchronize devices with satellite orbits. The results demonstrate a reduction in computational cost by up to 66.53% compared to state-of-the-art methods, with an accuracy comparable to more complex algorithms. Furthermore, the orbital information update mechanism ensures successful synchronization in over 87% of cases, significantly enhancing the energy efficiency and operational viability of DtS-IoT systems. This work bridges theoretical modeling and practical implementation, optimizing satellite-ground communication for constrained IoT deployments.

B.3.2.4 Parallelizing Orbital Propagators

The paper in [FFM13b], led by myself after receiving my engineering degree, tackles the inefficiency of sequential computation in satellite orbital propagation using the SGP4 model, a widely used algorithm for predicting satellite positions. Traditional implementations of SGP4 fail to scale for prominent satellite constellations due to their computational intensity, especially when real-time updates are required for mission-critical operations. The study introduces a parallel implementation of the SGP4 algorithm us-

ing OpenCL, a technique I contributed with in [FFM13a], leveraging GPU and multi-core CPU architectures to accelerate computation. The approach processes satellite orbital elements in parallel by exploiting data-level parallelism, efficiently distributing workload across hardware threads. The implementation is tested on various platforms, including GPUs and OpenCL-compatible CPUs. It optimizes kernel execution to balance computation and memory transfer overhead. The parallel SGP4 implementation achieves significant speed-ups, with performance improvements of up to 20x compared to traditional CPU-only methods, depending on the hardware configuration. The approach demonstrates scalability, maintaining high throughput as the number of satellites increases. These results underline the feasibility of real-time orbital propagation for prominent satellite constellations, enabling advancements in satellite operations, collision avoidance, and mission planning.

Appendix C

Routing and Flow Control

Routing and transport protocols are critical to ensuring the efficient operation of modern satellite networks, from DTN to LEO constellations. This section addresses two critical questions: *How to determine the next hop in multi-path satellite networks?* and *How to manage data flows effectively using advanced transport techniques?*

My work spans these domains, offering novel solutions to their unique challenges. In this section, I present my contributions, structured into two sections: **Low-Earth Orbit Constellations**, which highlights recent advances in routing and transport for highly dynamic satellite networks, and **Delay-Tolerant Networks**, where my foundational work during my PhD provided critical insights into time-aware routing and transport optimization for intermittently connected networks [FFB17a]. This progression reflects a natural evolution in my career, transitioning from DTNs, inspired by CONAE's SARE mission context, to scalable solutions for LEO mega-constellations with thousands of satellites. In both domains, my contributions emphasize the interplay between theoretical advancements and their practical application, capitalizing on a deep understanding of temporal and spatial dynamics in satellite communication systems.

C.1 Low-Earth Orbit Constellations

LEO constellations represent a paradigm shift in global connectivity, enabling unprecedented coverage and data rates through densely packed satellite networks. We covered their relevance in Section B.1 (see Figure B.1). However, these systems also introduce unique challenges, including dynamic topologies, frequent link changes, and stringent latency requirements, demanding innovative solutions at multiple layers of the communication stack. Besides the analysis summarized in Section B.1, my contributions to this field focus on specific methods in two critical aspects: **Routing**, where efficient and scalable algorithms ensure reliable data delivery, and **Transport**, where advanced

protocols adapt to the complexities of space networks to optimize throughput and latency.

C.1.1 Routing

Routing in LEO constellations is a cornerstone of efficient and scalable satellite network operations, as these constellations introduce unique challenges due to their dynamic topologies, rapid link changes, and integration with terrestrial systems. This section highlights two groundbreaking advancements in routing for LEO constellations: **SATNET-OSPF**, a protocol extension bridging terrestrial and space networks, and **DiscoRoute**, a distributed algorithm tailored for mega-constellations.

C.1.1.1 OSPF for Satellite Networks

In collaboration with Hongcheng Yan, from the China Academy of Space Technology (CAST), we addressed the critical challenge of enabling dynamic routing in integrated space-terrestrial networks characterized by time-varying topologies due to LEO satellite motion and environmental constraints [Yan+24a]. We found that existing terrestrial routing protocols, such as OSPF, lack adaptability to predictable and unpredictable topology changes inherent to satellite networks. In contrast, space-specific solutions often lack compatibility with terrestrial systems. We introduced SATNET-OSPF, an extension of the OSPFv3 protocol tailored for space-terrestrial networks to bridge this gap. Key innovations include (**a**) a memory-efficient adjacency list topology model capturing dynamic link states, (**b**) a fast link state detection mechanism leveraging decoding lock flags for wireless links, (**c**) proactive routing and forwarding mechanisms to handle predictable link-down events, and (**d**) integration with space-grade hardware using SPARC V8 CPUs and radiation-hardened FPGAs. SATNET-OSPF was implemented on a flight-ready space router and evaluated in a testbed simulating real-world space-terrestrial conditions. The details of the on-board router and the testing framework is presented in Figure C.1. Results demonstrate significant improvements, including a 31.4% reduction in protocol overhead, a four-orders-of-magnitude improvement in link detection time (from 40s to 192ms), a 22-fold decrease in routing outage during predictable link-down events, and consistent performance in key metrics such as route computation and forwarding delays. This work effectively adapts terrestrial routing protocols for the unique demands of space-terrestrial integration, providing a robust, scalable, and backward-compatible solution that paves the way for autonomous and efficient operation of future satellite constellations.

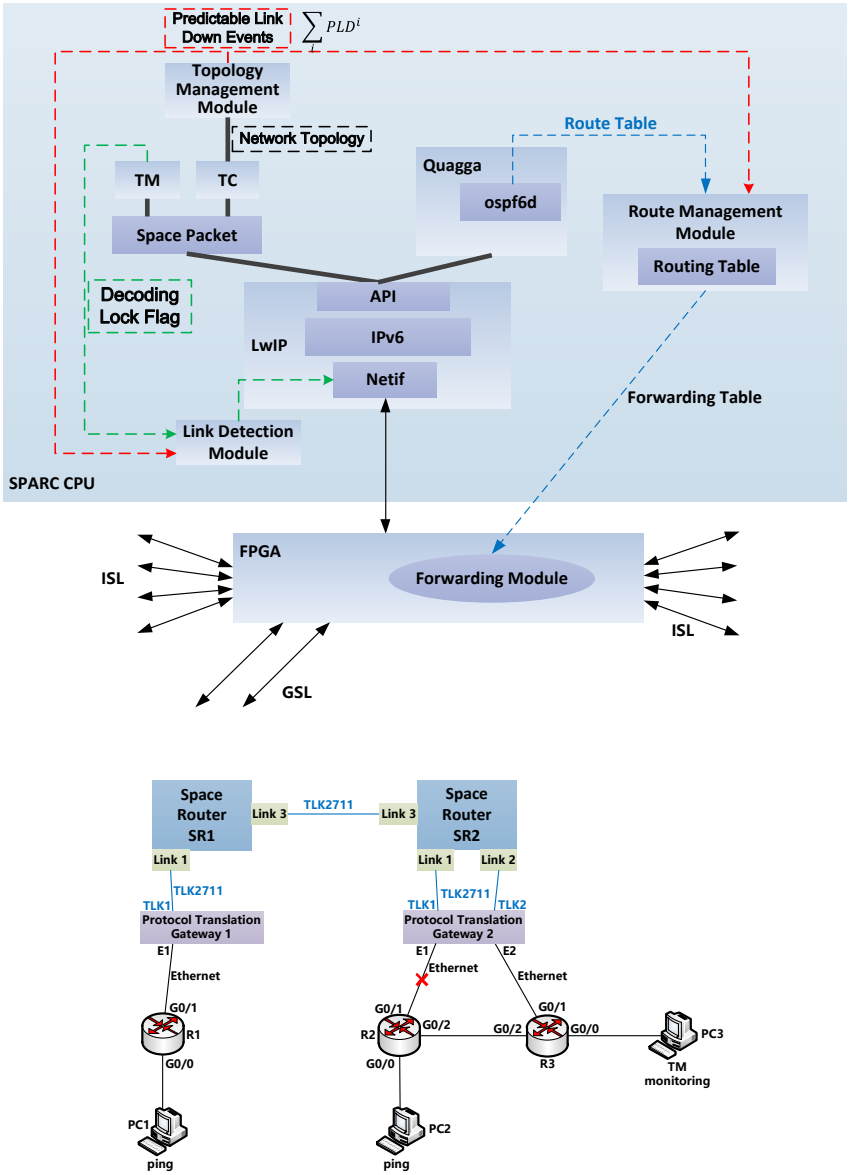

Figure C.1: SATNET-OSPF implementation and testing. **Top Panel**: Implementation details of the SATNET-OSPF router architecture. The SPARC CPU hosts the Topology Management Module (TM) for maintaining the dynamic network topology, the Link Detection Module leveraging Decoding Lock Flags to monitor link states, and Quagga's ospf6d daemon for OSPFv3 routing protocol management. The Route Management Module interacts with the Routing Table to update the Forwarding Table based on topology changes, including Predictable Link Down Events (PLD). The FPGA-based Forwarding Module ensures high-performance routing between Inter-Satellite Links (ISLs) and Ground-Satellite Links (GSLs), processing data at line rate. The architecture integrates lightweight IP (LwIP) and IPv6 stacks for efficient packet handling. **Bottom Panel**: Test setup validating space-terrestrial interactions by including classical terrestrial routers in the communication loop. Two Space Routers (SR1 and SR2), interconnected via ISLs and linked to terrestrial routers (R1, R2, R3) through GSLs, simulate operational conditions. TLK2711 protocol translation gateways bridge Ethernet connections to high-speed ISLs. (From [Yan+24a])

C.1.1.2 DiscoRoute

While SATNET-OSPF is focused on integrating with existing Internet routing solutions, such as OSPF, we took a different approach with DiscoRoute. With my supervised PhD student, Gregory Stock, we explored routing schemes that could take advantage of the orbital models to address the challenge of efficient routing in mega-constellations. The study in [SFH22a; SFH22b] develops a robust mathematical model to analyze satellite positions in Walker Delta constellations, deriving formulas for minimum hop counts between satellites. Leveraging these insights, the paper proposes two novel distributed, on-demand routing algorithms: a probabilistic algorithm and DisCoRoute. DisCoRoute minimizes hop count and travel distance with minimal computational overhead, exploiting patterns in inter-plane and intra-plane hop distances. These methods eliminate the need for exhaustive exploration or centralized pre-computation, ensuring scalability. Empirical evaluations focused on Starlink demonstrate that DisCoRoute is 158× faster than classical Dijkstra's algorithm and 7× faster than DAG-based methods while achieving near-optimal routing performance. Details of the orbital model and the results of DisCoRoute are presented in Figure C.2. The work marks a significant step forward in routing efficiency for mega-constellations. It presents solutions that scale effectively with constellation size and meet the demands of latency-sensitive space networks.

C.1.2 Flow Control

The evolution of satellite networks, particularly in the context of LEO constellations and mega-constellations, introduces unique challenges for flow control or transport protocols. High latency, elevated packet loss, and dynamic topologies stress congestion control algorithms, depending on the underlying routing strategies to optimize end-to-end performance.

Flow control in LEO Constellations are pivotal for enabling reliable, high-performance communication; however, my research reveals that their behavior in modern satellite networks' highly dynamic and latency-sensitive environment remains insufficiently understood. As one of the pioneering efforts in this domain, my work underscores the critical challenges and demonstrates the urgent need for a co-design approach that seamlessly integrates transport and routing layers to optimize performance and tackle the complexities of next-generation satellite communication systems.

My research in transport-layer optimization spans the **evaluation of BBR in QUIC** over satellite links and exploring the **impact of route stability in congestion control** per-

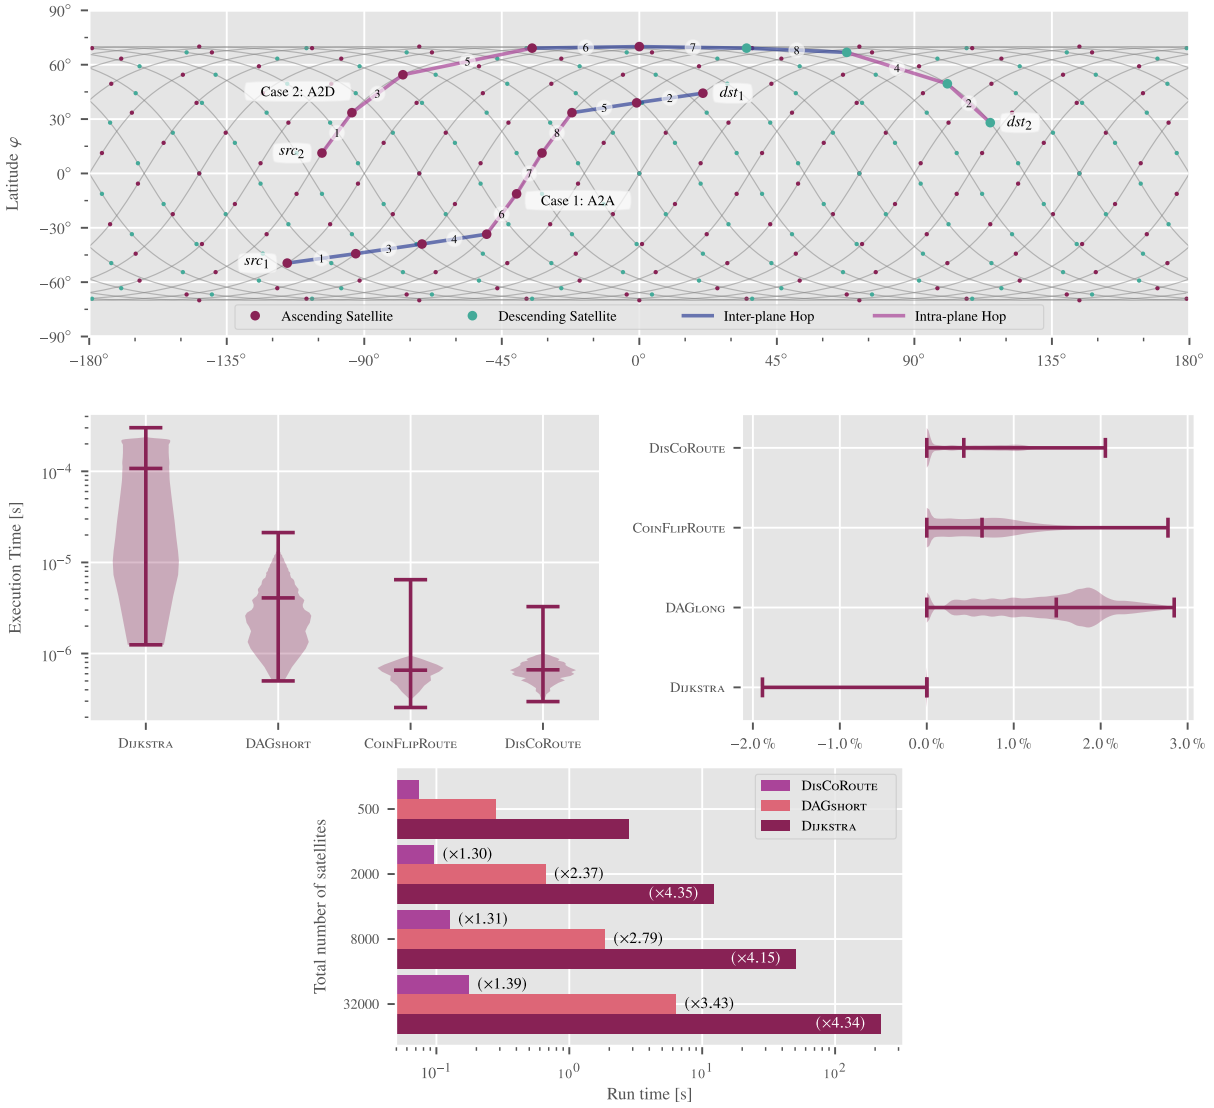

Figure C.2: Discoroute. **Top Panel**: Illustration of two routing cases that Discoroute discriminates in a Walker Delta constellation: Ascending-to-Ascending (A2A) and Ascending-to-Descending (A2D). A2A represents a routing scenario where hops occur between satellites within the same ascending orbit, while A2D involves hops between an ascending satellite and a descending satellite in adjacent orbital planes. **Center Left Panel**: Execution time comparison for various routing methods. Dijkstra's algorithm represents the classical approach, DAGshort leverages Directed Acyclic Graphs for pathfinding, and CoinFlipRoute employs a probabilistic heuristic. DisCoRoute outperforms all methods, demonstrating execution times orders of magnitude faster while maintaining competitive accuracy. **Center Right Panel**: The plot shows the percentage deviation of the computed paths from the shortest possible path. DisCoRoute exhibits minimal deviation, with a narrow distribution centered around 0%, showcasing its high accuracy. **Bottom Panel**: Scalability analysis for routing methods across constellations ranging from 500 to 32,000 satellites. DisCoRoute achieves exceptional scalability, with runtime consistently lower than alternatives, including DAGshort and Dijkstra. Even in the most prominent simulated constellations, DisCoRoute maintains efficiency, outperforming traditional methods by factors of up to 4×. (From [SFH22a; SFH22b])

formance in dynamic constellations. However, this is one of the most active research areas of writing, spanning lively cooperations with Nanjing University in China and the Saarland University in Germany.

C.1.2.1 Evaluation of BBR in QUIC

Our paper in [Wan+18] addresses the critical challenge of optimizing data transmission performance over satellite Internet, a domain inherently constrained by high latency and elevated packet loss rates. This work was led by Yue Wang, a Master's student I co-supervised with Prof. Kanglian Zhao from Nanjing University. The study evaluates the performance of Google's Quick UDP Internet Connections (QUIC) protocol integrated with Bottleneck Bandwidth and Round-trip propagation time (BBR), a congestion control algorithm designed to estimate bottleneck bandwidth and minimum round-trip time for high utilization without creating bottlenecks. We conducted extensive emulation-based evaluations using a GEO satellite Internet setup to achieve this. We tested under varying network conditions, such as high propagation delays (200ms to 1000ms) and packet loss rates (up to 20%). The findings reveal that QUIC with BBR consistently outperforms the classic Cubic Function to Congestion Control (CUBIC) congestion control algorithm, particularly in scenarios with high packet loss rates (>1%). While CUBIC's throughput and goodput degrade significantly under these conditions, QUIC with BBR demonstrates robust performance, maintaining higher throughput and stable goodput rates, making it well-suited for satellite Internet applications. Moreover, fairness evaluations indicate that BBR is compatible with CUBIC in multi-user bandwidth-sharing scenarios, though its probing behavior may temporarily reduce fairness. These results underline the effectiveness of QUIC with BBR in addressing the unique challenges of satellite networks. They also lay a foundation for further enhancements to congestion control algorithms tailored to satellite communications.

C.1.2.2 Congestion Control and Route Stability

While the BBR evaluation focused on ground-to-space and space-to-ground links, a complementary study in [Sto+24] explored the dynamics of ISLs in mega-constellations and their influence on congestion control algorithms. Gregory Stock, my PhD student from the Saarland University in Germany, led this work. In LEO constellations, many LEO satellites render highly dynamic topologies, resulting in frequent route changes that disrupt transport-layer protocols like TCP. Traditional routing methods, such as Dijkstra's algorithm, prioritize minimizing end-to-end delay but fail to consider the adverse effects of route instability, including packet reordering and unnecessary retransmissions. To address this, our paper proposes a suite of routing algorithms, including STUBBORN DI-

JKSTRA, TENACIOUS ROUTING, and SETCOVER, which focus on optimizing route stability while balancing end-to-end delay. The algorithms leverage predictable orbital dynamics of Walker Delta constellations to precompute routes that minimize instability. Using an evaluation framework called ARTCP, we simulated TCP performance over the proposed algorithms in Starlink-like constellations. Results show that TENACIOUS ROUTING achieves a 2x improvement in TCP data rate compared to Dijkstra by minimizing disruptive route changes. In contrastSETCOVER significantly reduces the number of changes without severely compromising performance. These algorithms demonstrate scalability, with runtime improvements up to 6x compared to traditional methods. They offer a practical solution for managing dynamic LEO networks and enhancing end-to-end communication reliability. The work lays the groundwork for co-designing routing and transport-layer protocols to optimize performance in satellite mega-constellations.

C.2 Delay-Tolerant Networks

This section explores routing and transport challenges in DTNs, which differ significantly from those in LEO constellations due to the critical role of time management. Unlike LEO networks, DTNs must accommodate data that can remain in transit for arbitrarily long periods over high-latency links or be stored onboard during extended link disruptions, such as planetary occlusions until connectivity is restored (see Figure C.4). As highlighted in my book [FFB17a], co-authored with Scott Burleigh, former NASA JPL and mentor of mine since my early career, introducing the time dimension into the networking domain presents exceptionally complex problems. These challenges far surpass those addressed by traditional protocols and algorithmic solutions designed for the terrestrial Internet, demanding innovative approaches to enable reliable and efficient communication in extreme environments.

> I find myself at a loss for words to express my gratitude to Scott Burleigh fully, retired from NASA JPL, and undoubtedly one of the most inspiring individuals I have encountered in my career. Beyond his remarkable technical expertise—he was the lead architect behind NASA's pioneering ION software—Scott introduced me to profound conceptual discussions about the critical role of time in space networks. Now a cherished friend, I am deeply thankful for every opportunity to meet and exchange ideas with him.

The remainder of this section explores DTN **routing with scheduled contact plans**, where future topology information is entirely accurate, **routing with uncertain contact plans**, where future topology information is uncertain or incomplete, and concludes with

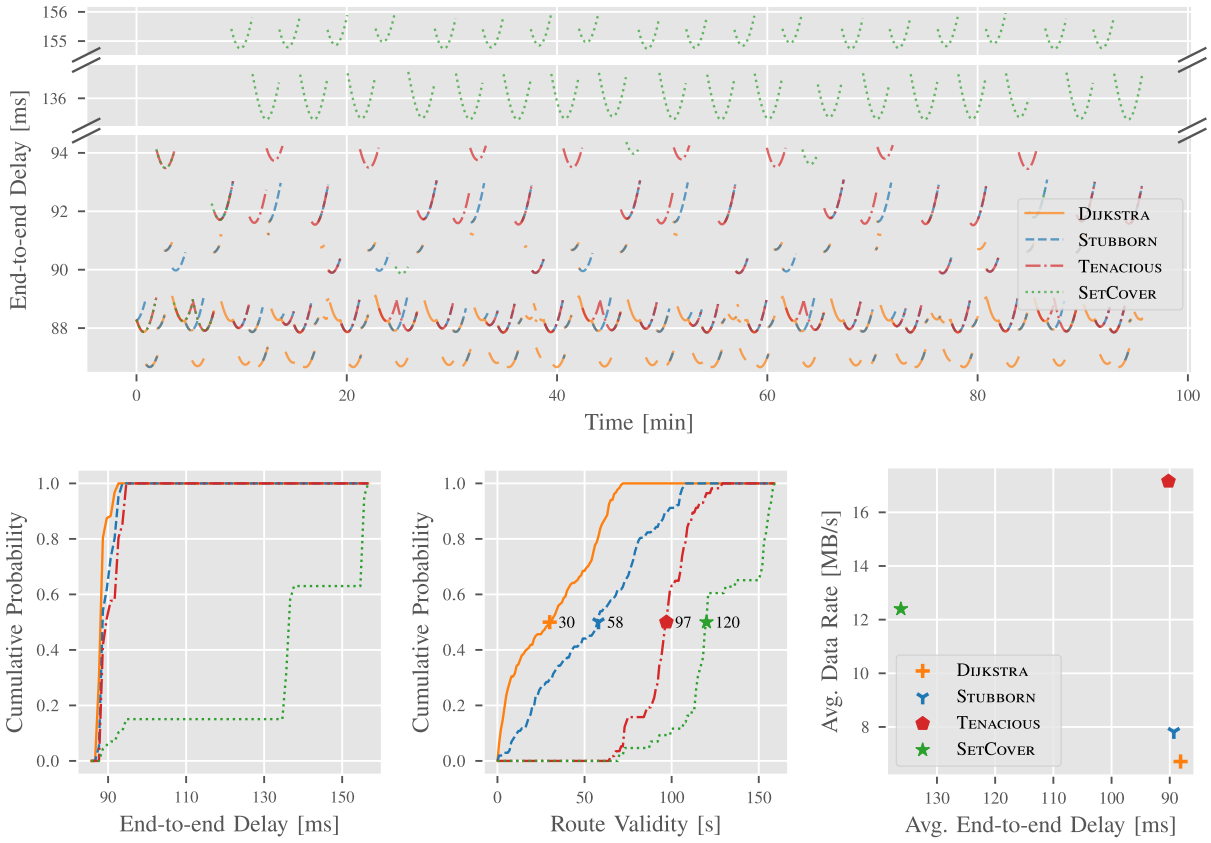

Figure C.3: **Top Panel**: The end-to-end delay dynamics across different routing schemes reveal latency and route stability trade-offs in a Starlink-like mega-constellation. Dijkstra achieves the lowest latency but requires frequent route changes, potentially complicating congestion control. In contrast, SetCover minimizes route changes, offering the most stable routing at the expense of significantly higher end-to-end delays. Tenacious and Stubborn routing approaches balance these extremes, with Tenacious particularly optimizing stability-latency trade-offs. **Bottom Left Panel**: Cumulative probability distributions of end-to-end delays highlight the latency performance of each routing scheme. Dijkstra dominates in latency but lacks route stability. **Bottom Center Panel**: Cumulative route validity times illustrate stability differences. SetCover maximizes route stability, while Dijkstra struggles to maintain consistent paths. Tenacious offers a middle ground with extended route validity. **Bottom Right Panel**: The average data rate plot evaluates transport protocol performance under routing-induced conditions. The results confirm that an optimal balance between latency and stability, as the Tenacious routing scheme approximates, enhances overall data throughput while mitigating congestion-related inefficiencies. (From [Sto+24])

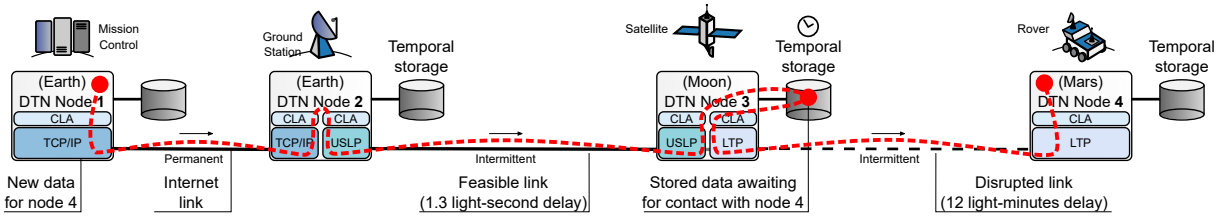

	Mercury	Venus	Earth	Mars	Jupiter	Saturn	Uranus	Neptune
Avg. Distance to Sun (AU)	0,38	0.72	1.00	1.52	5.21	9.54	19.18	30.11
Latency	3 min	6 min	8 min	**12 min**	43 min	1h20	2h40	4h10

Figure C.4: Store, carry, and forward example in a space DTN. A node on Earth transmits data to a satellite orbiting the Moon, which stores the data until a long-range contact with Mars becomes available. Because of a) delays and b) disruptions, reliance on immediate feedback messages must be reduced or eliminated from DTNs. Since a light-second is defined as the distance that electromagnetic radiation travels through free space in one second, it is a unit of distance that corresponds trivially to a unit of time. (From [FJB21])

an examination of **transport** considerations in DTNs.

C.2.1 Routing with Scheduled Contact Plans

Space assets benefit from highly predictable trajectories, which can be accurately forecasted using propagation models. These models account for various space effects, including planetary gravitational fields, atmospheric drag, third-body gravitational influences, solar radiation pressure, and other perturbations. Such models apply to both near-Earth and deep-space nodes. Similarly, whether stationary (e.g., ground stations) or mobile (e.g., rovers or drones), ground assets have predictable positions and visibilities that can be determined weeks in advance. This predictability enables the generation of a *contact plan*, which encapsulates the forthcoming network topology, as discussed in Section B.2. This capability is instrumental in designing and optimizing routing strategies in DTN space networks.

This section focuses on **Contact Graph Routing (CGR)**, arguably the most advanced, high-performing, yet complex routing approach for scheduled contact plans. Despite its strengths, CGR faces several open challenges in reliability, applicability, and scalability, where I have made significant contributions throughout my career. The section concludes with an introduction to **Shortest-Path Tree Routing for Space Networks (SPSN)**, a novel approach aimed at redefining the foundations of CGR by leveraging shortest-path tree algorithms to enhance its efficiency and adaptability.

C.2.1.1 Contact Graph Routing

Our paper in [FJB21], led by myself, presents a comprehensive tutorial on the problem of routing in the unique and challenging conditions of space DTN, where traditional terrestrial routing protocols are inadequate due to intermittent connectivity, long delays, and disruptions. The tutorial focuses on Contact Graph Routing (CGR), a framework tailored to leverage space networks' predictable yet dynamic topology derived from contact plans. To tackle these challenges, the paper revises the most relevant techniques to model time-varying DTN topologies, summarized in Figure C.5.

The paper proposes and elaborates on advanced routing methods for navigating contact-graph models, specifically focusing on the CGR framework. This framework encompasses algorithmic techniques that leverage contact graph models and optimized routing algorithms, including tailored Dijkstra's algorithm and Yen's K-shortest paths. Further discussed in Figure C.6, both variants enable efficient data planning, management, and forwarding in dynamic and delay-prone space networks. The tutorial was accompanied of CGR prototype code for direct evaluation of the approaches[1]. This code was integrated into NASA's ION (C code) and HDTN (Python interfacing) protocol stacks.

> I proposed adapting Yen's algorithm to compute the K-best shortest paths in a contact graph, addressing key inaccuracies in the legacy anchor approach, as detailed in [FJB21]. This enhanced variant improves DTN path selection accuracy and robustness. My adapted algorithm has been integrated into NASA's operational Interplanetary Overlay Network (ION) and High-Speed DTN (HDTN) software, currently deployed on multiple spacecraft, including the International Space Station (ISS). However, as I discuss in Section C.2.1.5, scalability aspects must be considered.

The key outcomes include an in-depth tutorial on CGR's operational stages—planning, routing, and forwarding—and the introduction of Python-based tools for simulation and implementation. **The CGR tutorial marks a landmark in my research path, consolidating decades of research and development into a unified methodology**. It demonstrates CGR's scalability, efficiency, and practical application in flight-grade systems like NASA's ION protocol stack, establishing it as the de facto routing framework for the evolving Space Internet.

Route Tables My paper in [Fra+18a] addressed the computational scalability challenges in route table construction for CGR. In this work, I first critique the existing *ini-*

[1]PyCGR public repository: https://bitbucket.org/juanfraire/pycgr/src/master/

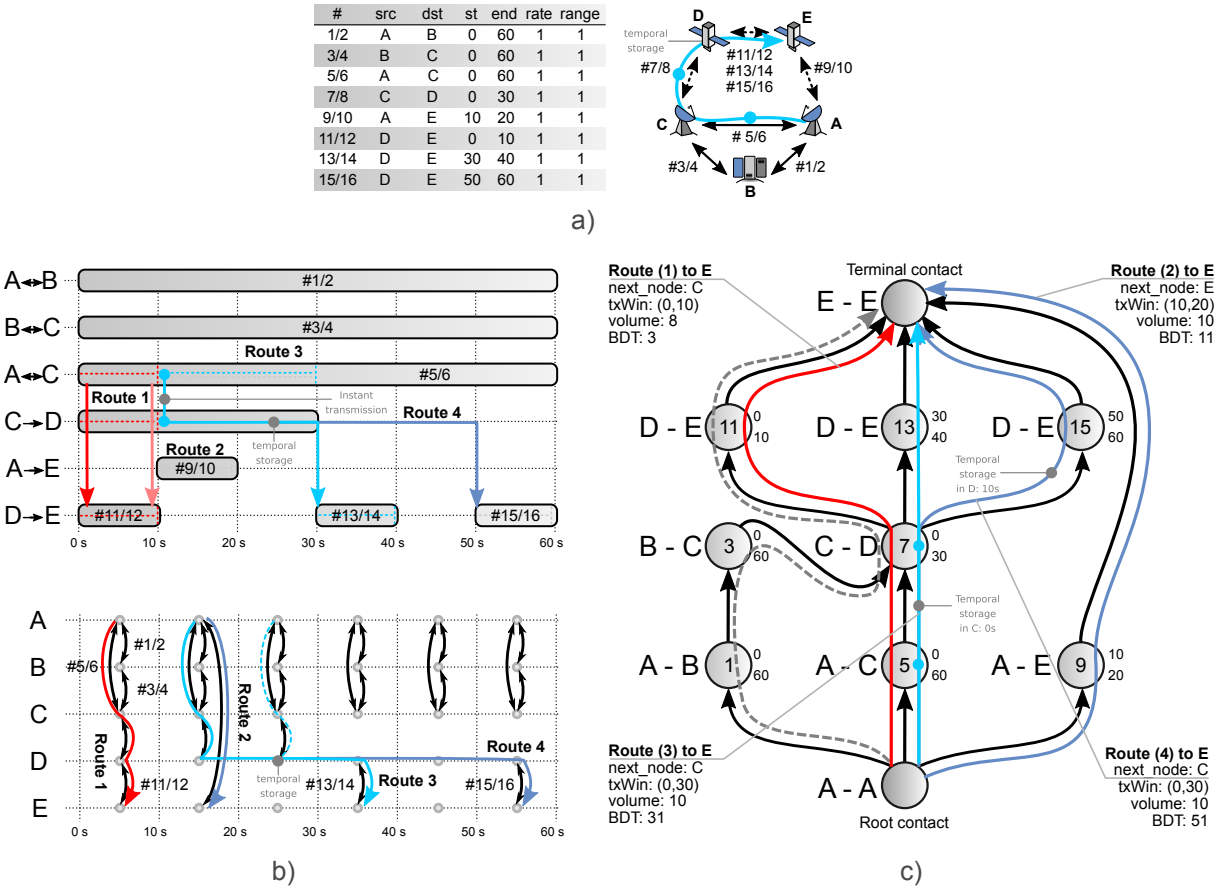

Figure C.5: Modeling DTNs using scheduled contact plans. **(a) Physical DTN Topology and Contact Plan**: Depicts the physical DTN topology with its corresponding contact plan, listing all possible contacts between nodes. Each contact includes details such as source, destination, start time, end time, data rate, and range, highlighting the temporal and spatial dynamics of the network. **(b) Classical Models**: The Timeline View (top) represents the contact plan as a time-ordered sequence of events, showcasing temporal connectivity between nodes. Example routes from node A to node E are highlighted. The Time-Evolving Graph (TEG) (bottom) visualizes the time-varying topology as a graph where edges correspond to available contacts over time. Routes from A to E are marked and annotated. **(c) Contact Graph Model**: Represents the same DTN contact plan as a compact Contact Graph where nodes and edges incorporate temporal and spatial connectivity information. The main feature of contact graphs is that they represent a time-evolving network in static graphs. Example routes from A to E are highlighted, ordered by BDT, with detailed annotations for transmission windows, volumes, and temporal storage durations. This model simplifies the decision-making process using simple adaptations to classical algorithms like Dijkstra for routing while maintaining accuracy, offering a robust framework for optimal data delivery in DTNs. (From [FJB21])

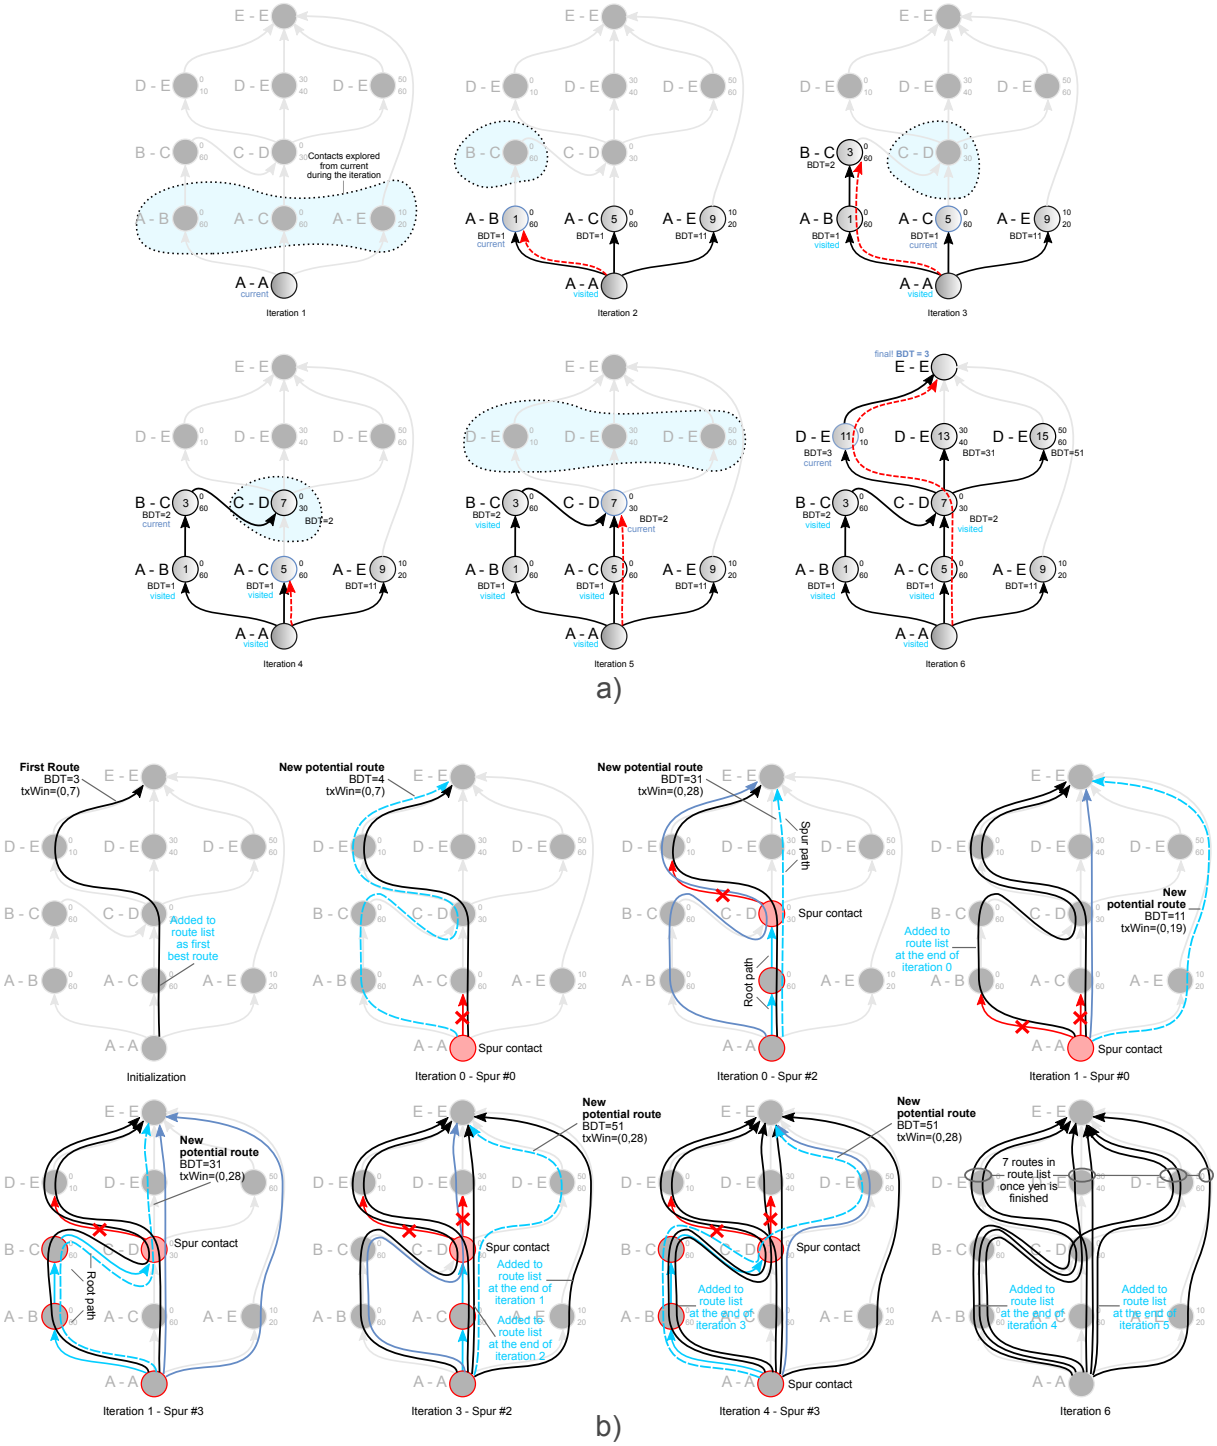

Figure C.6: CGR methods. **(a) Dijkstra's Algorithm Adaptation in CGR**: The Dijkstra-based CGR algorithm is shown iteratively identifying the shortest path from node A to node E based on the Best Delivery Time (BDT). The algorithm explores feasible contacts at each iteration, updating the BDT and selecting the optimal route while incorporating temporal constraints. The process terminates when the final path with the best BDT is found. **(b) Yen's Algorithm Adaptation in CGR**: Yen's K-shortest paths algorithm, adapted for CGR, iteratively constructs new candidate routes by avoiding previously traversed segments. Each potential route is evaluated for BDT and added to the list of K-shortest paths if it meets temporal feasibility and delivery performance criteria, enabling robust and redundant routing strategies in DTNs. (From [FJB21])

$R.BDT$	$R.EVL$	$R.tx_win$	Hops in R_A^E	Anchor	Ended	Depleted	Yens'	BFS
3	8	$(0, 10)$	$\{C_{A,C}^{0,60}, C_{C,D}^{0,30}, C_{D,E}^{0,10}\}$	✓	✓	✓	$K \geq 1$ ✓	✓
4	7	$(0, 10)$	$\{C_{A,B}^{0,60}, C_{B,C}^{0,60}, C_{C,D}^{0,30}, C_{D,E}^{0,10}\}$				$K \geq 2$ ✓	✓
11	10	$(10, 20)$	$\{C_{A,E}^{10,20}\}$	✓	✓	✓	$K \geq 3$ ✓	✓
31	10	$(0, 30)$	$\{C_{A,C}^{0,60}, C_{C,D}^{0,30}, C_{D,E}^{30,40}\}$		✓	✓	$K \geq 4$ ✓	✓
31	10	$(0, 30)$	$\{C_{A,B}^{0,60}, C_{B,C}^{0,60}, C_{C,D}^{0,30}, C_{D,E}^{30,40}\}$				$K \geq 5$ ✓	✓
51	10	$(0, 30)$	$\{C_{A,B}^{0,60}, C_{B,C}^{0,60}, C_{C,D}^{0,30}, C_{D,E}^{50,60}\}$				$K \geq 6$ ✓	✓
51	10	$(0, 30)$	$\{C_{A,C}^{0,60}, C_{C,D}^{0,30}, C_{D,E}^{40,50}\}$			✓	$K \geq 7$ ✓	✓

Figure C.7: CGR Route Table. This figure presents the list of routes discovered by various routing methods, including *anchor-based*, *first-ending*, *first-depleted*, *Yen's K-shortest paths* (with varying K), and *breadth-first search (BFS)*, applied to the contact plan topology shown in Figures C.5 and C.6. Each route is characterized by its **Best Delivery Time (R.BDT)**, **Expected Volume Left (R.EVL)**, and **Transmission Window (R.tx_win)**, alongside the respective hops within the contact graph. The *anchor-based* and *first-ending* methods identify basic routes but lack scalability for more complex scenarios. In contrast, *first-depleted* prioritizes routes with maximum throughput but can overlook certain paths when overlaps occur. As K increases in *Yen's algorithm*, progressively larger sets of routes are discovered. The *BFS* method iteratively explores all possible paths but lacks the route prioritization mechanisms inherent in algorithms like Yen's. (From [FJB21])

tial+anchor approach implemented in NASA's ION stack, highlighting its inefficiency in constructing large and static route tables that grow exponentially with topology size and traffic demands, especially in complex and dynamic DTSNs. I proposed innovative route table computation strategies to address these limitations, including *first-ending*, *first-depleted*, *one-route*, and *per-neighbor* methods, many of which are used in practical CGR implementations. Figure C.7 presents which route each of these techniques can be found among those present in the topology from Figures C.5 and C.6. These strategies dynamically update route tables on-demand, focusing on reducing computational overhead while maintaining forwarding efficiency. Simulation results over realistic scenarios demonstrate significant improvements in delivery ratios and computational effort. Notably, the *one-route* and *per-neighbor* approaches drastically minimize the size of route tables, avoiding unnecessary calculations and enabling real-time adaptability to topology changes. The study concludes that while no single strategy is universally optimal across diverse DTN configurations, these adaptive techniques markedly enhance scalability and operational efficiency compared to legacy methods.

C.2.1.2 Reliability

While CGR is positioned as the defacto routing routine for space DTNs, its reliability has been primarily unnoticed until my paper in [Fra+17b]. This work addresses the problem of assessing the performance and reliability of the CGR algorithm in distributed DTN satellite constellations operating under disruptive and high-latency conditions. To tackle this challenge, we developed a detailed framework combining the CGR algorithm with a robust fault model based on an exponential distribution, capturing the impact of transient node failures. The study introduces a new simulation tool, DtnSim, presented in detail in Section C.2.3. The setup efficiently analyzes scheduled DTNs, enabling the evaluation of CGR in realistic scenarios, including two satellite constellation formations: Walker Delta and Sun-Synchronous Along-Track.

The key outcomes revealed CGR's susceptibility to anchoring issues in complex constellations, highlighted significant rerouting overheads under failure conditions and demonstrated differences in fault resilience between the two formations. The Walker Delta formation exhibits superior overall performance but greater sensitivity to faults, while the Along-Track formation shows more stable performance with reduced sensitivity to increasing fault rates. These insights provided early advancements in understanding and optimizing CGR for operational DTNs. Routing techniques further expanded this work under uncertain contact plans, a topic we take on in Section C.2.2.

C.2.1.3 Applicability

Centralized and Decentralized Application Our early paper in [Mad+18a] and its journal extension in [FG21], partly work of Elias Gasparini, a Master student I supervised, delve into the applicability of centralized and decentralized routing approaches. The problem centers on the limitations of decentralized solutions like CGR, which demand significant on-board computational resources and autonomy—factors that can be risky for early space network deployments. To address this, I introduced two novel centralized routing schemes: Centralized first-ended (CFE) and Breadth-First Search (BFS), tailored to support both immediate forwarding (IF) and store-carry-forward (SCF) traffic (see Figure C.8). These centralized approaches offload computational demands to ground stations, leveraging superior processing capabilities while addressing space networking constraints like link instability and large-scale data transmissions. The study comprehensively evaluates CGR, CFE, and BFS and compares them using computational effort, memory utilization, and energy consumption metrics. It demonstrates that centralized methods are suitable for small—to medium-scale networks (up to 10k daily contacts) because they achieve computational efficiency and reduce energy costs. However, scalability challenges emerge for more extensive networks, reaffirming CGR's flex-

ibility for large-scale networks. The paper highlights the nuanced trade-offs between centralization and decentralization, advocating for hybrid approaches and further optimizations to bridge SIN routing requirements across diverse mission scales.

Application via Spot of Maximum Knowledge The paper in [FFW18] addresses the problem of applying CGR in a Lunar Ring Road Network, where low-cost low-lunar orbit (LLO) satellites serve as data mules to facilitate interaction between remote sensors, outposts, and Earth. The challenge lies in the time-varying network topology, exacerbated by lunar surface occlusions and the lack of continuous end-to-end paths. This complicates data's centralized tracking and routing. To address these issues, we proposed a novel routing strategy that utilizes a "Spot of Maximum Knowledge" (SMK) concept, where a subset of intensively interconnected nodes aggregates topological knowledge and scheduling decisions to optimize routing. This approach enables the central tracking of all communications by routing traffic through Earth mission control, supported by realistic simulation setups using STK and the Contact Plan Designer tool discussed in Section B.2.2.6. Key outcomes demonstrate that the SMK-based strategy achieves reliable routing and communication tracking while maintaining high delivery ratios and competitive average delivery times under various lunar network scenarios. The study underscores the feasibility of scalable, centralized lunar communication networks that can be deployed efficiently for future Moon exploration missions.

C.2.1.4 Viability

This part examines resource-efficient strategies for modeling, evaluating, and optimizing the utilization of contact volumes within the contact plan. While partially addressed in the other papers discussed in this section, our paper in [DFB24] addresses explicitly the inefficiencies in volume management for CGR, as defined in Consultative Committee for Space Data Systems (CCSDS)'s Schedule-Aware Bundle Routing (SABR) for DTNs, particularly the challenges of handling large bundles, overlapping contacts, and intermediate hop delays. It surveys current state-of-the-art approaches, such as the Effective Volume Limit (EVL), Earliest Transmission Opportunity (ETO), and queue delay (QD), to show that they fail to account for intricate routing scenarios, leading to underutilized resources and suboptimal delivery performance. Figure C.9 illustrates the volume management issues in interplanetary networks. To address these limitations, we proposed a novel contact segmentation technique that divides a contact into time-based segments. This technique enables precise tracking of available bandwidth and improves volume annotations for routing decisions. The paper introduces a generic volume management interface to enhance SABR's flexibility and modularity. This interface allows seamless integration of multiple volume management techniques. Evaluations

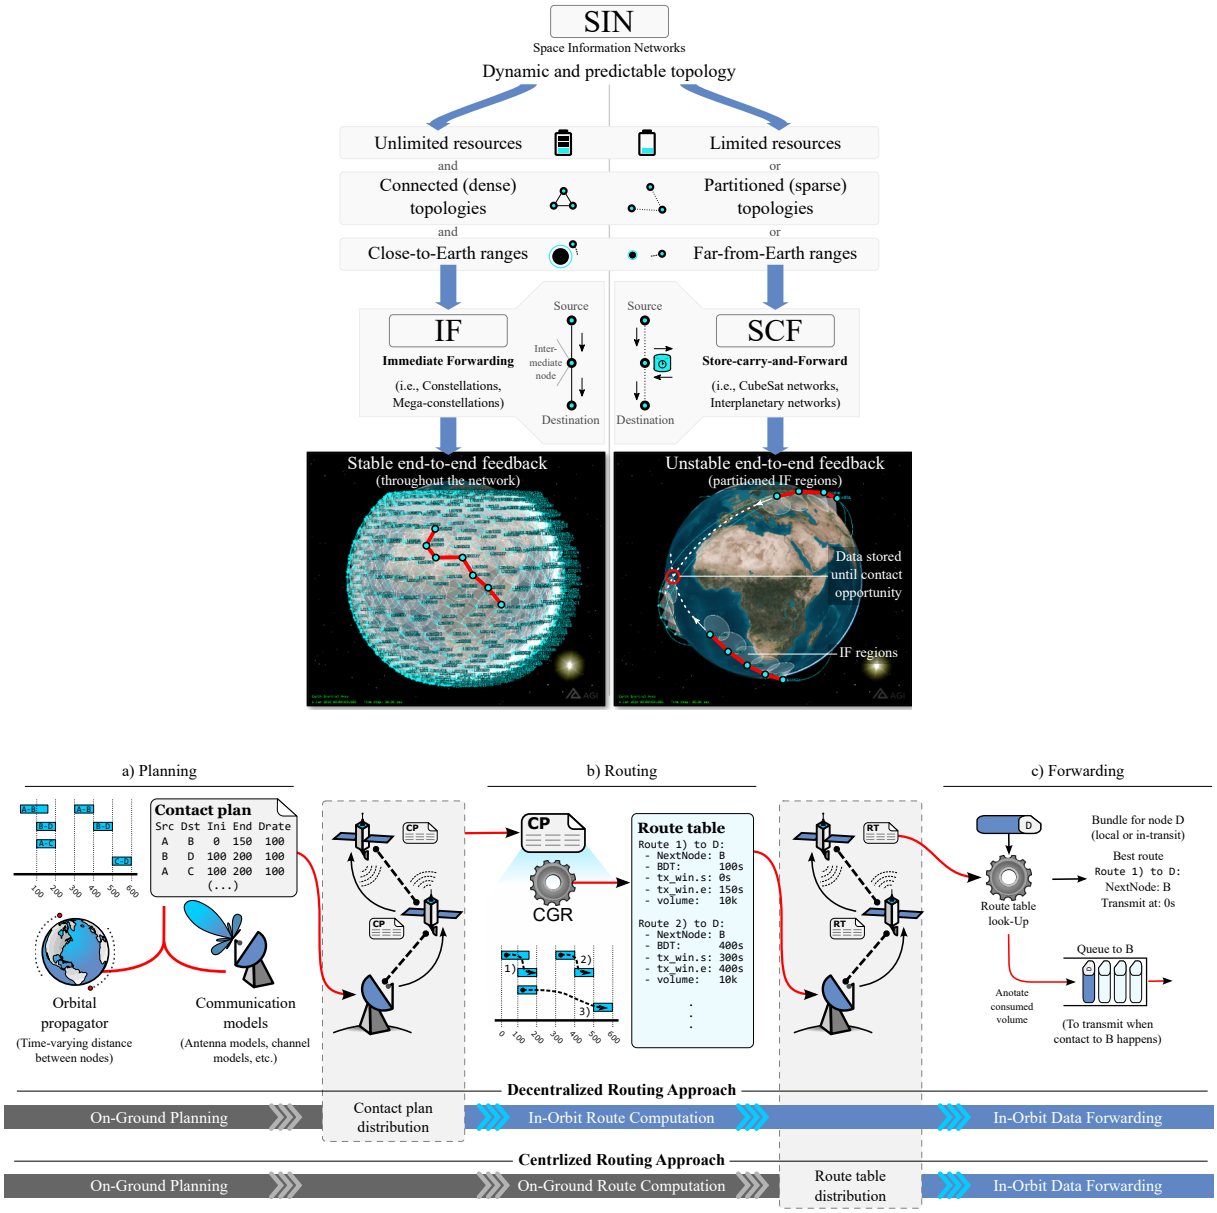

Figure C.8: Applicability of CGR. **Top Panel**: Space network topologies classification into dense, resource-rich networks like LEO mega-constellations (suited for immediate forwarding, or IF) and sparse, resource-constrained systems like CubeSat or interplanetary networks (requiring store-carry-and-forward, or SCF). While IF benefits from stable end-to-end feedback for real-time communications, SCF accommodates partitioned regions with intermittent connectivity by storing data until the next contact opportunity arises. **Bottom Panel**: Routing process phases: (a) Planning: Orbital propagators and communication models predict node positions, forming a contact plan that specifies connectivity windows. (b) Routing: CGR derives optimal paths using decentralized (in-orbit) or centralized (on-ground) approaches. Decentralized routing utilizes onboard computations for local traffic demands, while centralized routing precomputes and distributes route tables for the entire network. (c) Forwarding: Nodes consult route tables to place data bundles in outbound queues, ensuring timely transmissions when links become available. (From [FG21])

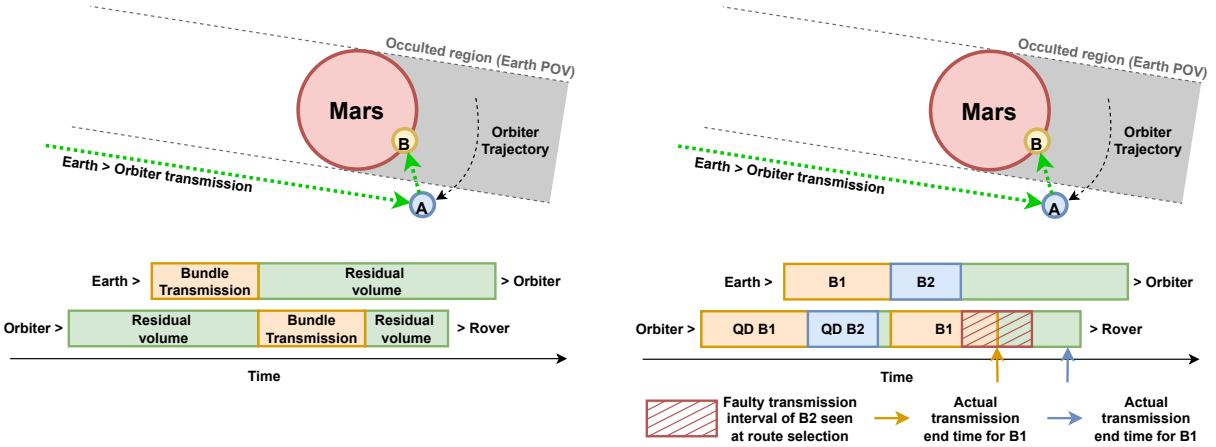

Figure C.9: Volume Management in CGR. **Left Panel**: The diagram shows a Mars relay orbiter communicating with Earth and a surface rover, highlighting the occulted region due to Mars' position from Earth's point of view. A sequence of bundle transmissions (e.g., B_1) is depicted, with residual volume carried forward as data is transmitted between Earth, the orbiter, and the rover. **Right Panel**: The timeline details bundles B_1 and B_2 scheduling, illustrating queue-delay (QD) effects. The faulty transmission interval of B_2, seen during route selection, results from the earlier start of an overlapping contact window. Despite B_1 inducing a queue delay ($QD\ B_1$), the actual transmission end time of B_1 prevents B_2 from starting as initially planned. (From [DFB24])

using realistic Ring Road Network (RRN) and Public Transport Network (PTN) scenarios demonstrate that contact segmentation achieves up to 10% higher delivery ratios, reduces end-to-end delays by up to 11%, and decreases hop counts by up to 16%, all without increasing computational overhead. These results establish contact segmentation as a robust solution for efficient volume management in scheduled space DTNs.

C.2.1.5 Scalability

Scalability in DTN routing is a double-edged sword. On the one hand, my contribution to introducing Yen's K-shortest paths algorithm in CGR marked a significant leap forward by enabling robust path computation for complex contact plans. However, this advancement also introduced a scalability bottleneck, as the iterative nature of Yen's algorithm leads to exponential growth in computational overhead and memory usage as network density and contact plan size increase. Practical solutions lie between these extremes—balancing comprehensive pathfinding with feasible onboard processing. The remainder of this section introduces my contributions along **scalability limits** analysis, **mitigation via onboard AI**, and **mitigation via regionalization**. I finally introduce the exploration of novel approaches called **SPSN**, which represent the path forward to address these challenges while maintaining operational efficiency in large-scale DTNs.

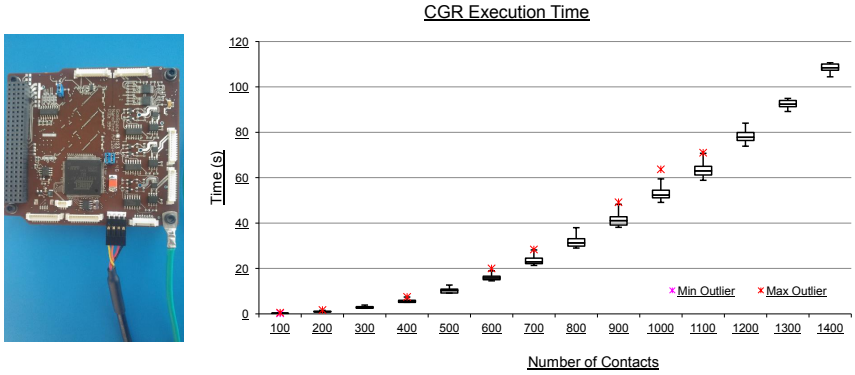

Figure C.10: CGR Implementation on GomSpace NanoMind A712C. **Left Panel**: The GomSpace NanoMind A712C, a flight-proven nanosatellite onboard computer (OBC) featuring a 32-bit ARM7 RISC CPU, 8 MB of SRAM, and limited computational capabilities, serves as the hardware platform for testing CGR's scalability. **Right Panel**: Execution time of the CGR algorithm as a function of the number of contacts in the routing plan. For small-scale contact plans (e.g., 100-300 contacts), execution times remain within feasible ranges (<10 seconds). However, as the number of contacts approaches and exceeds 900, execution times increase sharply, surpassing 60 seconds for 1,000 contacts and exceeding 100 seconds for 1,400 contacts. Such prolonged computation times are impractical for real-time satellite operations, underscoring the critical need for algorithmic optimizations to ensure scalability and mission feasibility in large-scale dynamic satellite networks. (From [VF20])

Empirical Scalability Limits The paper in [VF20], led by Blas Vega, a Master student I supervised from CONAE, addresses the scalability challenges of implementing CGR on nanosatellite onboard computers (OBCs) for dynamic DTNs with constrained computational resources. CGR's practical applicability on resource-constrained nanosatellites had not been empirically tested at scale before this paper. We implemented a streamlined version of CGR in C, optimized for the GomSpace NanoMind A712C OBC, which features a 32-bit ARM7 RISC CPU and limited memory capacity. We conducted extensive tests with contact plans ranging from 10 to 1,400 contacts to evaluate memory usage, execution time, and route discovery efficiency.

The results, summarized in Figure C.10, reveal that CGR exhibits polynomial scalability in terms of execution time, but memory limitations of the OBC impose practical constraints. For contact plans exceeding 900 contacts, memory exhaustion and increased execution times (up to 120 seconds) significantly hinder performance. Despite these challenges, the study demonstrates that CGR can operate effectively within the computational limits of nanosatellite OBCs for smaller-scale networks, paving the way for its deployment in realistic nanosatellite constellations while highlighting the need for further optimization in memory management and route computation.

The practical results found in the Nanomind OBC were further backed up by more theoretical analysis on Yen's algorithm scalability in [JFB23b]. The core problem ad-

dressed is the computational inefficiency and unpredictability introduced by Yen's K-shortest path algorithm when applied adaptively in CGR for large and dense networks, where routing table sizes and the number of iterations grow excessively, leading to memory and processing pressure. These inefficiencies are particularly pronounced when dealing with dense network segments and late-starting critical contacts, creating operational bottlenecks that degrade system performance. With Olivier De Jonckere, we analyzed CGR's scalability in-depth by evaluating different implementations and configurations, including contact and node parenting, and alternative strategies such as limiting-contact approaches. We simulated realistic satellite network scenarios (e.g., a ring road configuration) with 30 satellites and 30 ground stations to measure runtime overhead, delivery ratios, and processing pressures. Mitigation techniques, such as capping the maximum number of paths (K) and limiting Yen's routing iterations, are implemented to manage computational burdens. The results reveal that while limiting-contact strategies mitigate processing overhead, Yen's algorithm still imposes significant scalability constraints, particularly in high-density network scenarios. For instance, the computational overhead for Yen-based CGR grows by a factor of 98 when scaling from a 15-node to a 30-node scenario. Moreover, memory and processing pressures remain critical, as thousands of unnecessary routes are computed and discarded during scheduling. These findings underscore the need for alternative or hybrid routing strategies to ensure CGR's scalability in future large-scale delay-tolerant space networks.

Mitigation via Onboard AI The paper in [Olm+23], led by Matias Olmedo, my Master's student at CONAE, addresses the computational scalability limitations of CGR in delay-tolerant space networks. As discussed, as the number of nodes and contacts increases in space communication systems, the routing computation effort for CGR grows exponentially, posing challenges for onboard systems with constrained processing resources. To address this, the paper introduces GAUSS, a Graph Neural Network (GNN)-based routing approach designed to reduce computation time while maintaining reasonable network performance. GAUSS leverages supervised learning to pre-train a model on the ground using realistic contact plans. Then, it deploys the model onboard to predict routing decisions efficiently. The method utilizes an encode-process-decode pipeline, where GNNs inherently exploit the graph-structured data of space networks, ensuring scalability and stability in route computation. Evaluation results demonstrate that GAUSS halves the computational effort of CGR and reduces the variability in compute time by one-third across both near-Earth and cislunar scenarios. This makes GAUSS an appealing alternative for missions where minimizing onboard processing effort and ensuring consistent performance are critical. However, there is a trade-off in slightly suboptimal delivery time compared to CGR. Future work proposes fine-tuning GAUSS

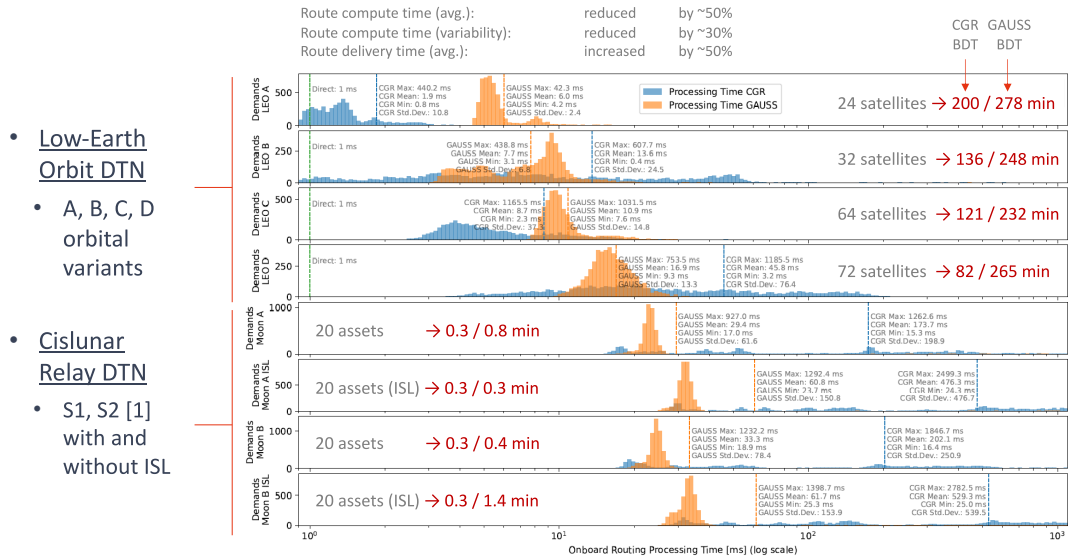

Figure C.11: Performance evaluation of the GAUSS algorithm. **Top Panel (LEO)**: GAUSS consistently reduces computation time by 50% compared to CGR over different LEO configurations, with lower variability (reduced by 30%). However, GAUSS increases delivery time (by 50%) due to its approximate routing approach. **Bottom Panel (Cislunar)**: GAUSS achieves faster onboard routing with minimal computational variability in cislunar cases. Key results show GAUSS reducing the maximum computation time to 0.3 minutes compared to CGR's 0.8–1.4 minutes. (From [Olm+23])

to improve delivery metrics further while retaining its scalability benefits.

Mitigation via Regionalization Our paper in [Mad+18b] was led by Pablo Madoery during his PhD I help supervise and tackle the scalability challenge of routing in large-scale scheduled DTNs handling vast contacts and nodes. We introduced two regionalized knowledge mechanisms—Global Regionalized Knowledge (GRK) and Local Regionalized Knowledge (LRK)—that aim to constrain the topology knowledge available to each node, effectively partitioning the network into manageable regions connected by passageway nodes. GRK reduces the contact plan size by removing inter-region contacts unless they involve passageway nodes, while LRK further localizes routing by assigning region-specific contact plans to nodes, relying on hierarchical routing with an inter-region routing algorithm (Inter-RR) for inter-regional communication. The definition of such an Inter-RR algorithm is the core contribution of this paper. Simulations demonstrate that these regionalized approaches significantly improve scalability, with LRK offering the most significant computational savings by limiting nodes' routing computations to local regions. However, the trade-off is slightly reduced in network performance, particularly in LRK, as fewer routes are available and static inter-region links are used instead of time-varying ones. The results highlight the flexibility of regionalized routing mechanisms in balancing scalability and performance, providing a scalable framework

for DTN routing in large-scale space networks.

Our paper in [JF23] challenges the model *nodes as passageway* model presented in [Mad+18b]. Instead, this paper, led by Olivier De Jonckere, proposes a novel *contact as passageway* model that abstracts regional borders as virtual nodes using graph contraction techniques. By merging multiple inter-regional contacts into single virtual vertices, this approach reduces the size of regional graph representations while maintaining routing flexibility and simplifying pathfinding operations. The proposed method introduces the concept of regional shortcuts, transforming the traditional regional tree structure into a graph-based topology to enhance connectivity and performance. Evaluations demonstrate that this approach significantly improves delivery ratios, with gains of up to 11% in complex scenarios, reduces delivery delays by over 17%, and lowers transmission overhead by more than 14%.

Last but not least, with Felix Flentge, we supported the supervision of Nicole Wenzinger from the ETH Zurich, entitled *Scoping Routing Scalability in Space Delay- and Disruption-Tolerant Networks*. The manuscript, available online since 2023, does an excellent job at summarizing and comparing the methods in the state of the art of DTN topology regionalization, including our two papers mentioned above.

C.2.1.6 Shortest-Path Tree Routing

The paper in [DF20], led by Olivier De Jonckere, addresses the scalability challenges of CGR in DTNs, mainly when dealing with large-scale contact plans with thousands of contacts. The proposed solution, SPSN, revises the underlying graph model of CGR, replacing the contact-graph paradigm with a more traditional nodes-as-vertices structure. This adjustment enables the adoption of efficient graph algorithms, such as spanning tree searches and priority queues, which are computationally infeasible under CGR's original model. SPSN calculates routes to all destination nodes in a single execution instead of CGR's iterative and source-destination-specific process. This early evaluation demonstrates that SPSN reduces computation times by more than an order of magnitude compared to CGR, especially for large-scale scenarios, making it a practical alternative for modern space networks. These results underscore the potential for SPSN to replace legacy CGR in scenarios demanding high scalability and computational efficiency while maintaining route optimality.

A follow-up paper in [JFB23a] also led by Olivier De Jonckere further deepened on SPSN. We introduced two key innovations: *multipath-tracking* and *tree-caching*, embedded in the SPSN framework. Multipath-tracking mitigates the limitations of traditional Dijkstra-based approaches by simultaneously exploring multiple valid paths, thereby reducing the need for iterative recalculations. Tree-caching extends this by enabling the reuse of previously computed shortest-path trees for subsequent routing decisions, sig-

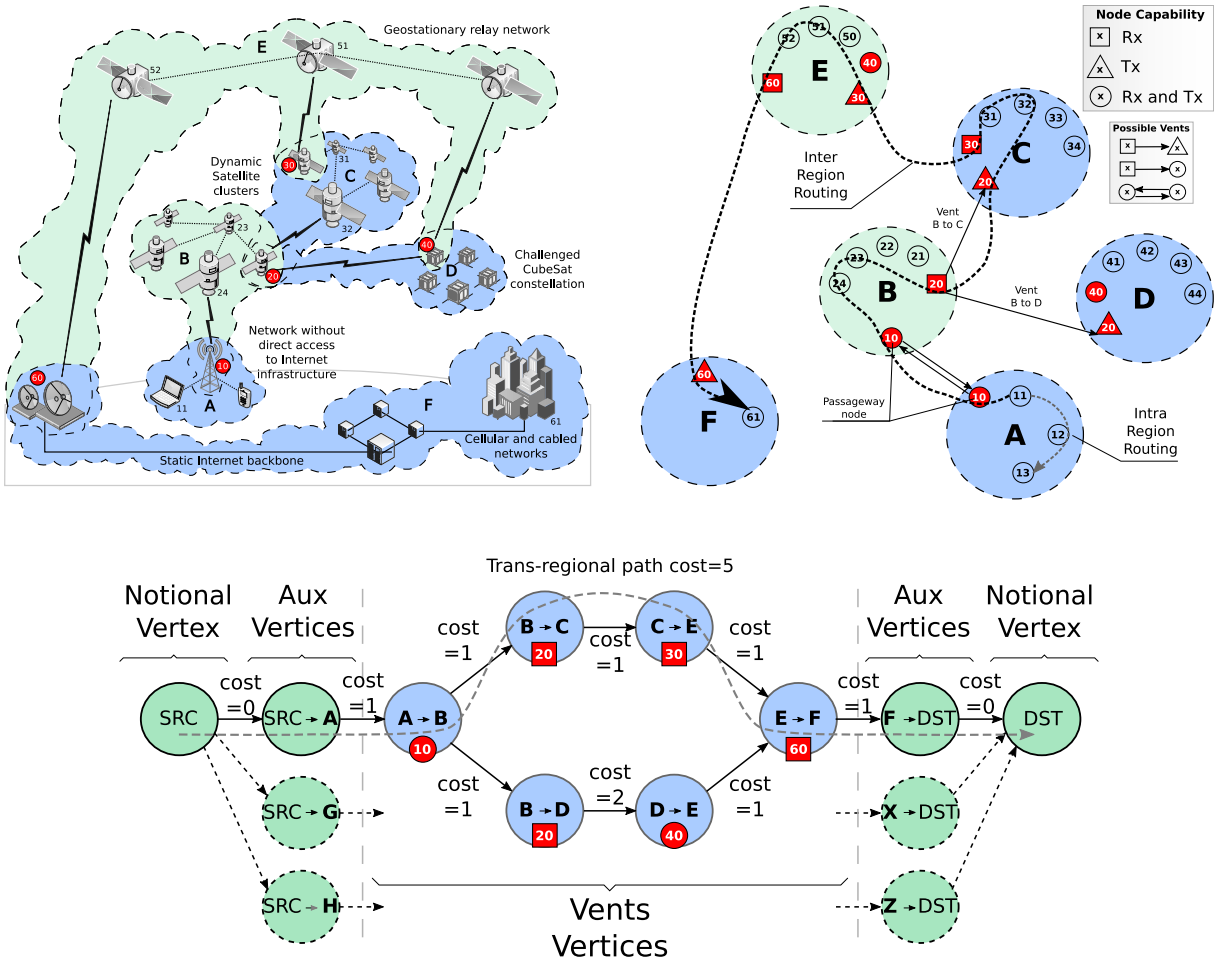

Figure C.12: Regionalized Routing in Large-Scale DTNs. **Top Left Panel**: Illustration of region splitting in a large-scale space DTN. The network is divided into regions (A-F) based on geographical or operational boundaries, such as dynamic satellite clusters, static internet backbones, or challenged CubeSat constellations. **Top Right Panel**: Depiction of the regionalized topology. Regions are interconnected through *vents* (*passageway nodes*) that enable inter-region communication. Intra-region routing is handled independently within each region, utilizing localized contact plans. **Bottom Panel**: Illustration of the Inter-Regional Routing (Inter-RR) mechanism. Auxiliary vertices represent passageway nodes, while notional vertices represent the source and destination nodes in a regionalized contact graph. For instance, a trans-regional route from A to F is computed by combining intra-region paths and inter-region links through passageway nodes (vents), yielding a total path cost of 5. The Inter-RR algorithm computes the trans-regional path cost using the shortest path approach across auxiliary vertices. (From [Mad+18b])

nificantly reducing redundant computations. Simulation results reveal that SPSN minimizes the computational load and sustains delivery rates under high-traffic scenarios. Regarding runtime efficiency and scalability, it outperforms classical CGR implementations. This work positions SPSN as a scalable and practical alternative for large-scale space DTNs.

Olivier De Jonckere completed his PhD at the Technische Universität Dresden, with SPSN as the centerpiece of his research. During this time, I had the privilege of informally co-supervising his work. Despite the notable advances and promising preliminary results, SPSN remains conceptually challenging to convey. As of this writing, we continue to collaborate with Olivier from his current position with the MAORE team in Montpellier, striving to enhance SPSN's visibility and impact.

C.2.2 Routing with Uncertain Contact Plans

CGR and its variations, as introduced in Section C.2.1, operate under the assumption of perfectly accurate deterministic contact plans, where every predicted contact occurs as planned. However, this assumption rarely holds in real-world scenarios, even space networks, where factors such as antenna misalignments, interference, and transient faults can disrupt otherwise predictable links. In such cases, CGR exhibits reduced performance, necessitating robust approaches to address the inherent uncertainties in contact plans.

We first introduced the routing concept under uncertain contact plans in the context of Fernando Raverta 's PhD thesis, which I co-supervised alongside Prof. Jorge M. Finochietto. This pioneering effort investigated the challenges of addressing uncertainty in contact plan predictions, leading to the exploration of replication-based schemes over contact plans.

Uncertain contact plans refer to scheduled communication opportunities where predicted contacts' availability, timing, or quality may deviate from initial forecasts due to various unpredictable factors such as environmental interference, hardware malfunctions, or resource constraints. Routing in uncertain contact plans entails developing algorithms that optimize multi-copy forwarding decisions under probabilistic and incomplete network knowledge. These algorithms leverage modeling, replication, and learning techniques to balance delivery reliability, latency, and resource efficiency.

This section follows a chronological progression of Fernando Raverta 's PhD work, beginning with **CGR-2Copies**, a straightforward dual-copy extension of CGR. It then ad-

vances to a formal problem framework using **Markov Decision Processes** and concludes with scalable, localized solutions leveraging **statistical** and **learning-based approaches**.

C.2.2.1 CGR-2Copies

Our papers in [Mad+18c] and [Rav+18b], led by Pablo Madoery and Fernando Raverta respectively, were the first to introduce the challenge of routing in DTNs under uncertain contact plans, where scheduled connectivity may deviate due to transient faults, interference, or other unpredictable factors. The paper introduces alternative routing metrics and a novel replication-based CGR variant called CGR-2Copies to address this. This method combines the optimal delivery-time route with a minimum-hop count path to balance speed and reliability. Simulations are conducted in random and realistic satellite network scenarios, including a Walker constellation pattern, under varying levels of contact uncertainty. Results, summarized in Figure C.13, show that CGR-2Copies significantly improve delivery ratios and energy efficiency compared to standard CGR, particularly at moderate uncertainty levels while maintaining acceptable delay and hop metrics. These findings highlight the potential of combining deterministic and opportunistic routing strategies to enhance DTN performance under non-ideal conditions. This was the stepping stone work that showed the tradeoff between delivery time and reliability as delivery metrics in DTNs.

C.2.2.2 Markov Decision Processes

The work with CGR-2Copies highlighted the pressing need for a more formal framework to address routing under uncertain contact plans. To tackle this, we collaborated with Prof. Pedro D'Argenio from FAMAF (UNC), whose expertise guided us to the domain of Markov Decision Process (MDP). This paradigm proved to be an ideal fit for modeling and solving this complex problem.

The first approach to MDPs, led by Fernando Raverta, was the paper in [Rav+18a]. This work introduced a theoretical framework for optimal routing decisions when contact failures occur. To address this, we modeled the routing problem using a MDP implemented in PRISM Model Checker, creating the so-called Best Routing Under Failures (BRUF) model. BRUF leverages probabilistic transitions to encode contact failure probabilities and computes optimal routing policies that maximize delivery success. Experimental evaluations in simulated DTNs with varying contact failure probabilities demonstrate that BRUF is an upper theoretical bound for delivery ratios in single-copy routing. Compared to CGR and its variants, BRUF achieves superior reliability under increasing uncertainty, while CGR with replication strategies (e.g., CGR-2Copies) improves delivery

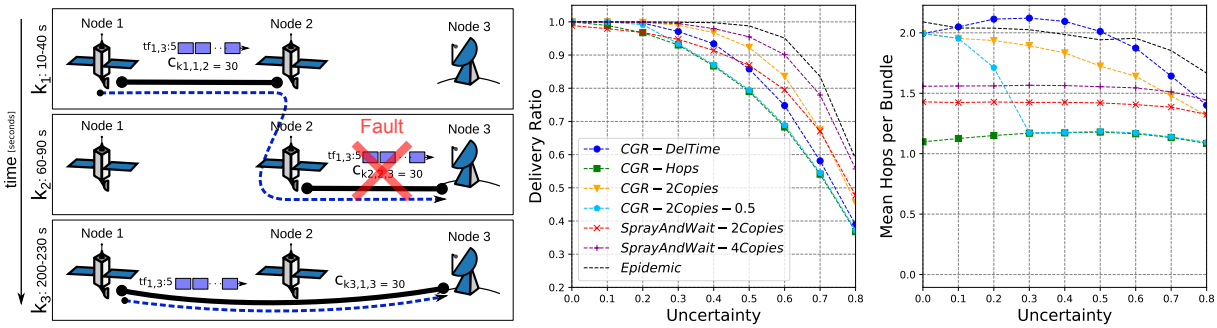

Figure C.13: CGR-2Copies. **Left Panel**: An illustration of a contact fault in a DTN, demonstrating a disrupted traffic flow due to an unexpected fault in contact k_2. The optimal flow would have rerouted through contact k_3 instead, highlighting the need for fault-tolerant routing strategies. **Center Panel**: Delivery ratio comparison across various routing methods under increasing contact uncertainty. Standard CGR variants (based on Delivery Time and Hop metrics) suffer significantly as uncertainty grows. The CGR-2Copies approach shows robust performance, balancing reliability and efficiency, outperforming standard CGR. While Spray-and-Wait (S&W) with 4 copies and Epidemic routing achieve higher delivery ratios, they do so at the cost of significantly higher energy consumption due to increased transmissions. **Right Panel**: The mean hop count, serving as a proxy for energy consumption, highlights the trade-off between delivery ratio and resource efficiency. CGR-2Copies achieves a balanced performance with lower hop counts. (From [Mad+18c])

ratios but at the cost of energy efficiency. The results also highlight the BRUF framework's scalability limitations for large networks due to the high computational overhead of constructing the decision matrix. This necessitated future research into scalable and approximate routing solutions for uncertain DTNs.

This last problem was addressed in the following paper [Rav+21], possibly the core contribution of Fernando Raverta in his PhD.. We introduced a novel theoretical framework, Routing under Uncertain Contact Plans (RUCoP). RUCoP leverages a hand-made Markov Decision Process MDP formulation to model uncertain contact plans and provide optimal routing decisions that maximize delivery probabilities. The method encodes all possible routing states and actions, including failure detection and rerouting mechanisms, into a structured algorithm that computes the delivery probabilities for multiple-copy scenarios. Key outcomes from the study demonstrate that RUCoP improves delivery ratios by up to 25% compared to CGR and Spray-and-Wait under realistic LEO satellite scenarios with uncertain inter-satellite and ground contacts. Practical adaptations, such as Local-RUCoP (L-RUCoP) and a CGR variant (CGR-UCoP), were proposed for deployment on existing DTN protocol stacks, showing that these methods closely approximate RUCoP's theoretical performance while ensuring practical feasibility in resource-constrained DTN nodes. These results position RUCoP and its derivations as robust solutions for dynamic and fault-prone space communication networks.

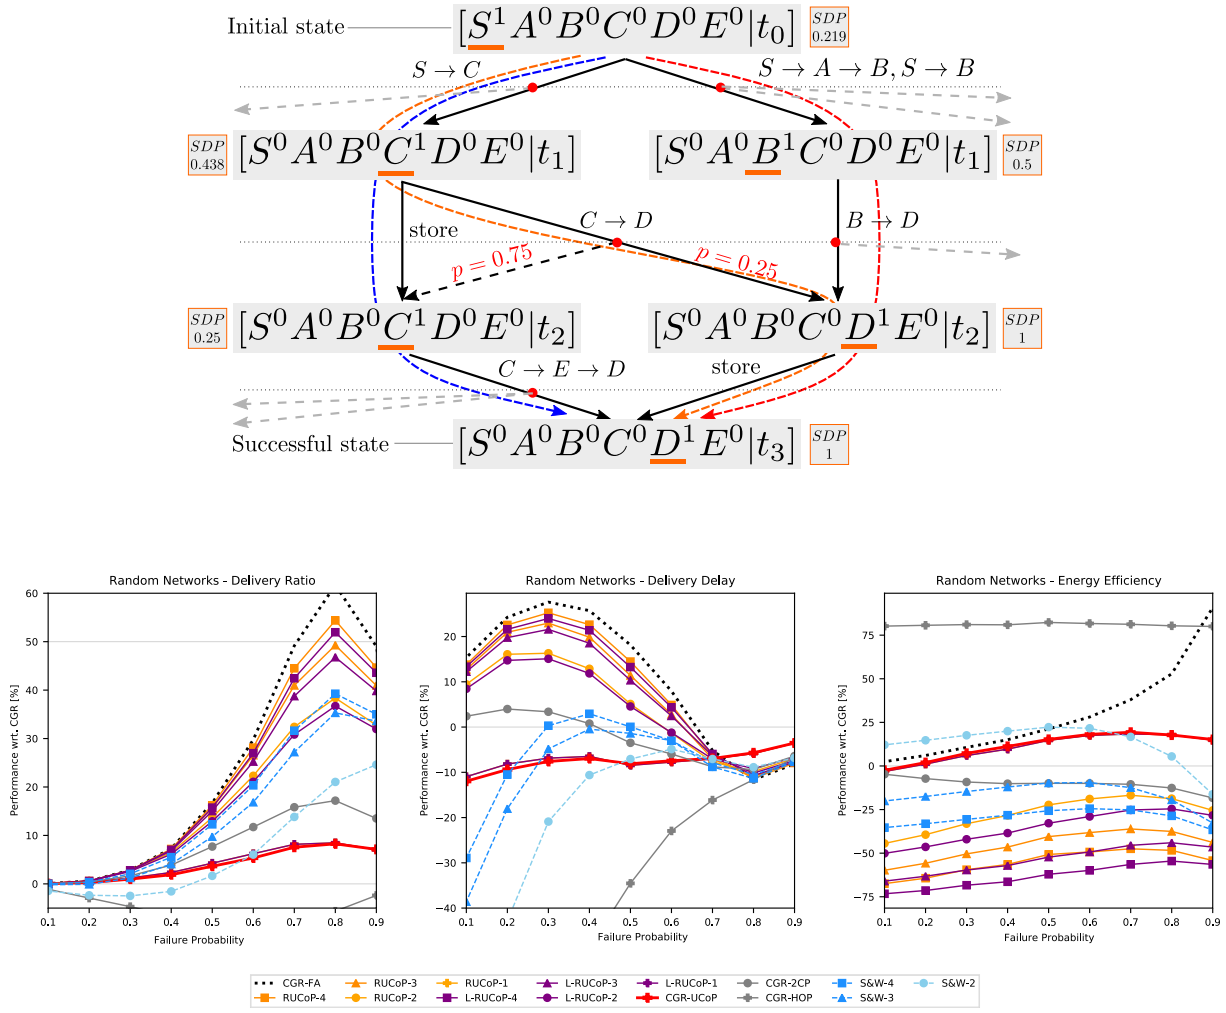

Figure C.14: RUCoP. **Top Panel**: The MDP state formulation illustrates the system's non-deterministic nature, where each state represents the configuration of bundle copies across nodes at a specific timestamp. Orange bars on nodes indicate the presence of a copy at the corresponding state. Transitions between states involve either transmitting the bundle to another node (depicted by arrows) or retaining the bundle in storage (store action). Transition probabilities (e.g., p = 0.25) capture the likelihood of successful transmissions or faults. The Successful Delivery Probability (SDP) is computed from the bottom (successful states) to the top (initial state). **Bottom Panel**: Performance evaluation of delivery ratio, delivery delay, and energy efficiency across various routing schemes in a random network generator. RUCoP variants with 1, 2, 3, and 4 copies consistently outperform others in delivery ratio and delay, with L-RUCoP (local RUCoP) following closely, providing comparable performance with reduced computational complexity. Energy efficiency metrics indicate that simpler schemes like Spray-and-Wait (S&W) are more efficient in energy usage but fail to deliver high delivery ratios or low delays. The results demonstrate RUCoP's superiority in balancing delivery performance and delay, while L-RUCoP offers a realistic alternative with minimal trade-offs. (From [Rav+21])

At this stage, we recognized that while RUCoP delivered optimal results and featured an implementation capable of offloading the state graph to disk, the method faced significant scalability challenges. Furthermore, the assumption of a global vision of the complete network status, as discussed in detail in Figure C.15, limits the practical usefulness of RUCoP. These limitations prompted us to explore alternative approaches to address the growing complexity effectively.

C.2.2.3 Statistical Approaches

Our paper in [DFH20], developed in collaboration with Prof. Pedro D'Argenio and the involvement of Fernando Raverta, analyzes alternative ways of exploiting the MDP model. Traditional global schedulers, computed using Probabilistic Model Checking (PMC), are unimplementable due to their reliance on global state information, as discussed in Figure C.15. This is the case of RUCoP and the core reason for the L-RUCoP and the CGR variants, which only rely on local information. The paper introduces a novel framework for distributed schedulers to address this, enabling local decision-making based solely on node-specific information. We adopted the Lightweight Scheduler Sampling (LSS) technique, initially introduced by Prof. Pedro D'Argenio, from Statistical Model Checking (SMC) to generate implementable routing strategies for these distributed schedulers. Using the Modest Toolset, we built a toolchain to model DTNs with realistic contact plans and evaluate the proposed approach on a Walker constellation case study. The results demonstrate that the distributed LSS method reliably finds high-probability routing strategies that are implementable in real-world DTNs, with comparable performance to global-information schedulers in terms of delivery probability while maintaining scalability and low computational overhead. This establishes LSS as a practical and efficient alternative for resilient DTN routing under uncertain conditions.

A more detailed evaluation of the role of statistical approaches for routing in uncertain DTNs was presented in [DAr+22], focusing on their scalability and solution quality trade-offs. The study benchmarks RUCoP, which employs an exhaustive analytical exploration of MDPs, and LSS, a simulation-driven statistical method. The evaluation reveals that while RUCoP achieves superior reliability in smaller, structured networks, LSS outperforms it in computational efficiency for larger, more complex topologies. Both methods exhibit limitations in scenarios with extensive uncertainty or high failure probabilities. These findings highlight the need for hybrid approaches to balance scalability and performance in real-world DTN applications.

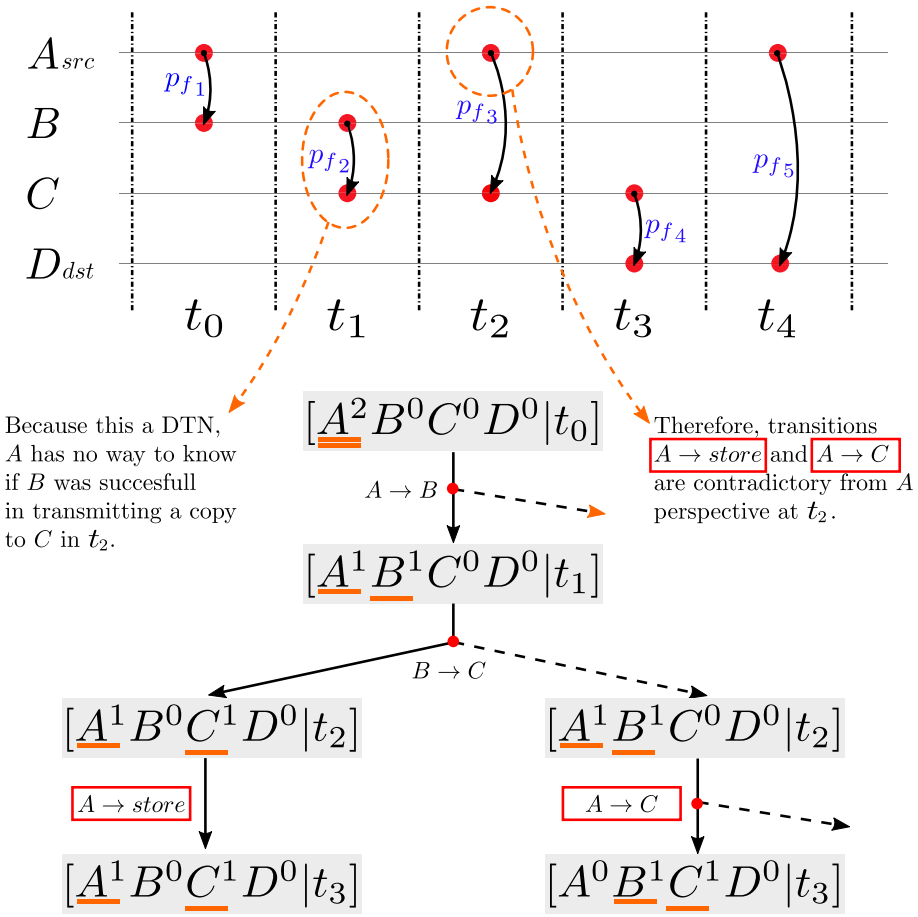

Figure C.15: Global vs. local information in RUCoP. In a DTN, local nodes (e.g., node A) cannot ascertain the current distribution of bundle copies across the network due to the lack of instantaneous feedback. For instance, at time t_2, node A cannot determine whether node B successfully forwarded the bundle to node C or retained it. This uncertainty results in contradictory transitions in the RUCoP MDP, where A might simultaneously assume both *store* and *forward* states, depending on global knowledge of the network that is unavailable in practice. The original RUCoP approach assumes the availability of global state information, simplifying the modeling but impractical in real-world DTN scenarios. To address this, the Local RUCoP (L-RUCoP) and the CGR variant (CGR-UCoP) adaptations introduce heuristics to approximate the routing decisions based solely on local information. Decision-making is restricted to what the local node can infer, such as the history of its transmissions and local contact probabilities, rather than relying on the complete state of the network. This adaptation reduces the reliance on unrealistic global visibility while providing reasonable routing performance in constrained and dynamic DTN environments. (From [Rav+21])

C.2.2.4 Learning Approaches

Our paper in [DAr+24] explored the role of learning approaches to address routing under uncertain contact plans. It introduces the application of Q-Learning (QL) focused on the scalability and adaptability issues posed by existing methods like RUCoP and LSS. Unlike RUCoP's exhaustive exploration of the MDP or LSS's reliance on random scheduler sampling, QL employs a Reinforcement Learning (RL) approach that iteratively updates a Q-table to approximate the quality of routing actions under uncertainty. This allows for more flexible exploration and exploitation of the network's state space. Furthermore, the paper proposes a localized variant, L-Q-Learning (QL) (L-QL), tailored for DTNs where nodes lack global information. By focusing on distributed learning, L-QL achieves a linear growth in memory usage relative to the number of nodes, compared to the exponential growth seen in global QL. The study benchmarks QL and L-QL against RUCoP and LSS using various topologies, showing that QL consistently delivers competitive delivery probabilities concerning LSS, particularly in simpler topologies and exhibits resource scalability advantages in complex networks. Figure C.16 presents a more detailed discussion of the scenarios and key outcomes. However, this work also reveals that excessive exploration episodes yield diminishing returns in performance gains, highlighting the efficiency of concurrent QL in distributed settings. These findings position QL, especially its local adaptation, as a viable middle-ground solution for scalable and efficient DTN routing.

C.2.3 Tooling: DtnSim

All routing results presented for scheduled contact plans (Section C.2.1) and uncertain contact plans (Section C.2.2) were achieved using DtnSim, our OMNeT++-based simulator[2]. Over the years, DtnSim has been progressively developed and refined, evolving into one of the most comprehensive and widely used libraries for DTN routing research and evaluation.

Our paper in [Fra+17a] introduced the core architecture of DtnSim, a discrete-event simulator designed to evaluate routing algorithms and network behavior in space-terrestrial DTN. The problem addressed is the lack of tools capable of seamlessly integrating flight-grade software (e.g., NASA's ION) with scalable simulation environments, particularly for large-scale, time-evolving topologies spanning near-Earth and interplanetary networks. The event-driven framework accelerates simulations while integrating actual flight software such as the CGR module from ION. This enables replicating operational routing decisions within a controllable simulation environment, making it suitable for validating and optimizing DTN algorithms. Key features of DtnSim include: (1) the ability

[2]DtnSim public repository: https://gitlab.inria.fr/jfraire/dtnsim

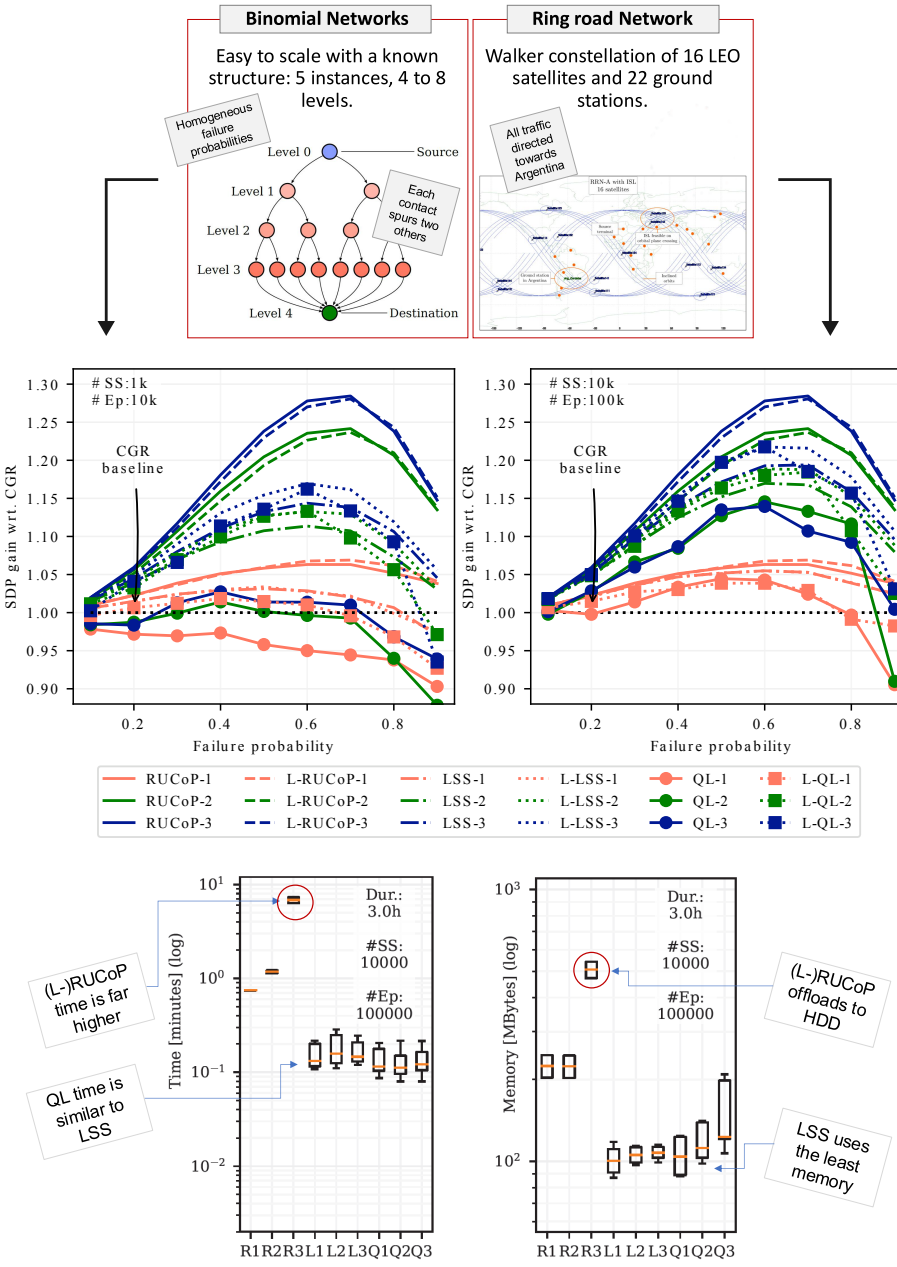

Figure C.16: Comparing RUCoP, LSS, and Q-Learning. **Top Panel**: Two evaluation scenarios are presented: the Binomial Networks, artificial topologies with a scalable structure designed to stress routing algorithms, and the Ring Road Network, a realistic setup based on a Walker constellation of 16 LEO satellites. **Center Panel**: The Successful Delivery Probability (SDP) gain is evaluated relative to CGR. The left figure represents low Scheduler Samplers (SS) for LSS and limited episodes for Q-Learning, whereas the right figure corresponds to larger SS and episode counts. Results indicate that RUCoP and L-RUCoP consistently outperform other schemes across most failure probabilities, while QL and LSS show competitive performance in specific failure regions. **Bottom Panel**: Computational time and memory resource usage for the evaluated schemes reveal a significant trade-off. While RUCoP and L-RUCoP deliver the best performance, their computational and memory requirements are considerably higher compared to LSS and QL. LSS demonstrates the lowest memory footprint, whereas QL balances resource consumption. (From [DAr+24])

to simulate complex DTN topologies over extended durations with high computational efficiency; (2) integration of realistic contact plans for precise evaluation of CGR and other routing protocols; and (3) support for rapid prototyping and educational use, with public availability for community adoption.

DtnSim was showcased in a recent survey about existing simulation and emulation platforms in cooperation with colleagues from NASA [Bul+24]. This paper addresses the critical challenge of creating realistic, high-fidelity network emulation environments to prototype and validate DTN protocols and software for space applications. To address this challenge, we evaluated various emulation platforms and tools, such as DtnSim, NASA's DTN Engineering Network (DEN), CORE/EMANE, and the Glenn Research Center's Automated Test Suite, emphasizing their capabilities to simulate real-world conditions like long delays, intermittent connectivity, and constrained resources. The study highlights how emulation enables accurate testing across multiple DTN layers, integrating tools like Netropy for delay and jitter modeling, Docker for containerized network node management, and STK for physical layer integration. Key outcomes include demonstrating the effectiveness of emulation in advancing DTN technologies by validating routing protocols (e.g., Epidemic and cognitive routing), measuring security impacts (e.g., BPSec), and preparing for mission-critical deployments in cislunar and ISS scenarios. It is worth mentioning that this paper also demonstrated the integration of IPN-v, our visualization tool presented in Section A.2.2.2, with DtnSim and other tools from NASA. We underscored the need for standardized emulation frameworks and scenario templates to ensure cross-platform compatibility and enhance collaboration between academic, governmental, and industrial stakeholders.

C.2.4 Flow Control

Flow control in DTNs presents unique challenges due to the influence of the temporal storage of data and high latency links, which necessitates a strict hop-by-hop layer design. While hop-by-hop flow control can be effective for short-delay, single-hop links by leveraging immediate feedback, this approach becomes increasingly inefficient under high-latency conditions, such as those encountered in interplanetary communications. High propagation delays disrupt the timely exchange of feedback, leading to inefficient data flow, suboptimal resource utilization, and the risk of network congestion. Congestion in DTNs also refers to storage congestion, not only contact overbooking.

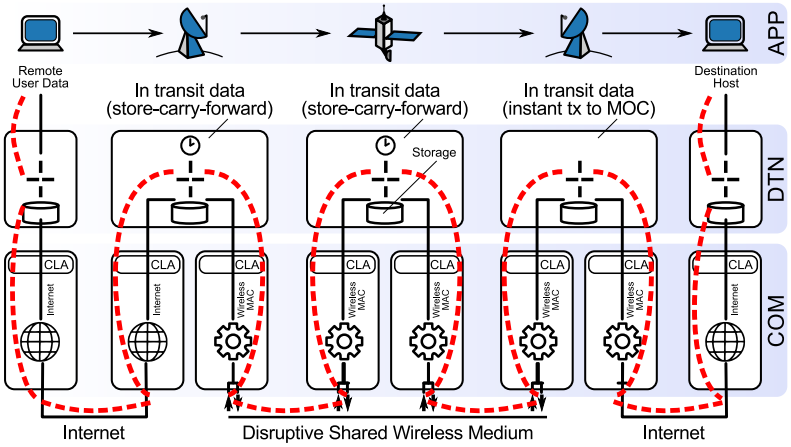

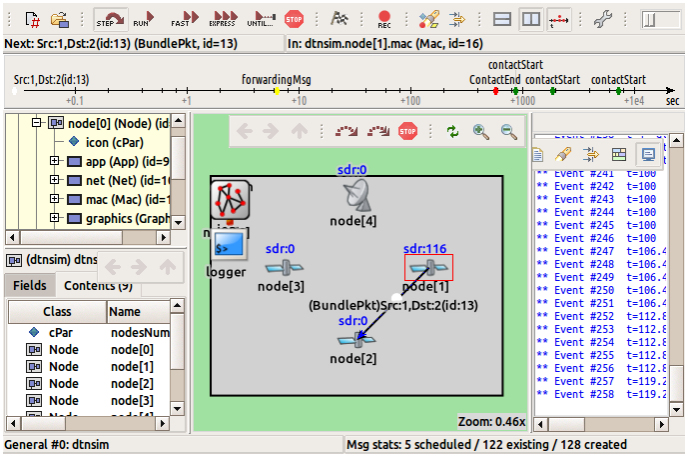

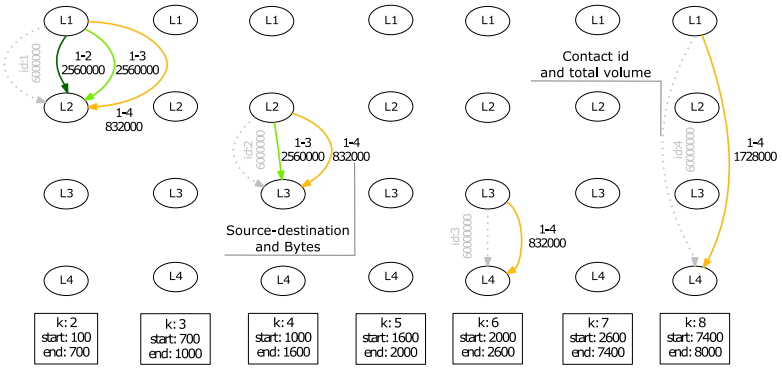

Figure C.17: DtnSim. **Top Panel**: DtnSim architecture highlights its three core modules: APP, representing the application layer that generates and processes data; DTN, which implements the store-carry-forward mechanism, interacting with the CLA for interoperability across heterogeneous protocols; and COM, encapsulating the communication stack, including wireless MAC and physical layers. **Center Panel**: A screenshot of DtnSim's GUI. This interactive GUI enables detailed debugging and analysis of DTN routing. **Bottom Panel**: A sample time-evolving topology output graph showcasing contact schedules and packet flows. Nodes represent network elements, while directed edges represent contact opportunities with annotated source-destination pairs and data volumes. (From [Fra+17a])

Flow control in DTNs is inherently complex due to the interplay of high-latency links and limited temporal storage, requiring advanced mechanisms to manage data flow efficiently. Unlike traditional networks, DTNs face contact overbooking and storage congestion, necessitating innovative strategies to ensure reliability and optimize resource utilization under extreme delay conditions.

These challenges demand advanced methods for managing flow control beyond traditional approaches, integrating mechanisms to address reliability and congestion across extreme delay environments. The remainder of this section details my contributions to flow control, structured into two key areas: **reliable delivery** and **congestion control**.

C.2.4.1 Reliable Delivery

Most of my contributions to reliability have been made in collaboration with Prof. Kanglian Zhao 's team at Nanjing University (China) and Prof. Ruhai Wang 's team at Lamar University (USA), both experts in the topic, particularly in Licklider Transmission Protocol (LTP), a transport-layer protocol designed for reliable data delivery over long-delay and lossy space communication links.

Quantifying File Delivery Time Our first paper in [Cao+18] addresses the challenge of quantifying the expected file delivery time of the BP in space networks characterized by asymmetric data and acknowledgment (ACK) channel rates and significant propagation delays (but no disruptions). To tackle this, our study develops an analytical model for BP's performance, incorporating parameters such as file size, bundle size, channel rates, retransmission intervals, and channel error rates. The methodology includes modeling retransmissions caused by corrupted or delayed ACKs and conducting experimental validation using a testbed running BP over emulated space channels. Key outcomes reveal that channel asymmetry severely impacts file delivery time, mainly when the ACK channel rate is lower than the data channel rate. The study finds that increasing the bundle size improves performance up to a point, but further increases can lead to diminishing returns due to retransmission inefficiencies. The proposed model closely matches experimental results, demonstrating its reliability for predicting BP performance in asymmetric and delay-prone environments. It also highlights the critical need for optimized retransmission strategies to mitigate delivery delays.

Impact of Disruptions Our paper in [Zho+19] addresses the impact of intermittent link disruptions on reliable file delivery using the BP and Licklider Transmission Protocol

(LTP) in a relay-based communication architecture between the ISS and ground stations. The analysis explicitly evaluates three disruption scenarios: (1) interruptions on the ISS-to-relay GEO satellite link, (2) interruptions on the relay-to-ISS return link affecting acknowledgment processes, and (3) disruptions on the relay-to-ground station link. The method involves analytical modeling to estimate file delivery times under varying disruption intervals and numerical evaluations considering parameters such as bit error rate (BER), retransmission timeouts (RTOs), and link asymmetries. Key findings demonstrate that more prolonged link disruption intervals significantly increase retransmission efforts and file delivery times. Under extreme conditions, the disruption effects overshadow channel errors. These insights underscore the critical need to optimize protocol parameters and relay configurations to enhance data delivery efficiency in intermittently connected space networks.

Energy-Efficient Delivery Our paper in [Liu+19] investigates the energy consumption of the BP for reliable file delivery over deep-space channels, comparing traditional *reactive* retransmission with *proactive* retransmission. The problem addressed is the significant energy demand for retransmissions in space communications, where energy resources are critically limited, particularly under lossy channel conditions and long propagation delays. We presented an analytical model to estimate total transmission efforts, considering bundle size, bit-error rates (BERs), and retransmission strategies. Numerical evaluations reveal that proactive retransmission, designed to resend data based on anticipated losses preemptively, consumes more energy than the traditional reactive method due to the repeated transmission of all bundles within a single round-trip time. However, proactive retransmission ensures improved transmission efficiency by reducing overall delays. Key findings highlight the trade-offs between energy efficiency and reliability. They emphasize that larger bundles and higher BERs exponentially increase energy costs, and proactive methods are particularly energy-intensive under adverse conditions.

C.2.4.2 Congestion Control

DTN congestion control has long been a challenging research topic. During my PhD, I was among the first to explore the intricacies of congestion management, authoring several pioneering papers as the lead researcher. Subsequent advancements were achieved under my co-supervision of Pablo Madoery from UNC, Argentina, and Lei Yang from Nanjing University, China.

Congestion Analysis My first paper on congestion [FMF14b] tackles the computational inefficiencies and lack of congestion avoidance in routing mechanisms for predictable DTNs, focusing on CGR and Merugu's Floyd-Warshall (MFW). This early DTN routing approach never transitioned to practical deployment. To address these issues, we proposed Cache-CGR (C-CGR), an enhanced version of CGR that leverages a caching mechanism to reduce computational overhead while preserving local congestion avoidance capabilities. C-CGR dynamically stores and reuses next-hop routing decisions for bundles within a contact, avoiding repetitive executions of CGR. The method is evaluated through simulations and a test bench using NASA's ION implementation, demonstrating significant reductions in computational load and stable throughput even in high-traffic scenarios. Key findings reveal that C-CGR achieves comparable routing performance to CGR while requiring fewer computations, making it particularly suitable for resource-constrained spaceborne DTN nodes. Moreover, the paper highlights the need for global congestion avoidance mechanisms as an open research challenge for scalable DTN systems.

Congestion Modeling My second paper [Fra+15a] deepened on the critical issue of congestion management in predictable DTNs, where traditional CGR mechanisms face limitations in handling system-wide congestion due to their focus on local contact capacities. To overcome these challenges, I proposed three solutions: (1) a Linear Programming (LP) model that establishes theoretical congestion-free routing and forwarding bounds for predictable topologies (this was a foundational model then used for contact plan design models discussed in Section B.2.2), (2) Path-Aware CGR (PA-CGR), which extends CGR to account for the full path capacity, and (3) Multi-Graph CGR (MG-CGR), which generates node-specific contact plans based on traffic predictability and optimizes resource usage while avoiding congestion. These methods are rigorously evaluated using a realistic LEO satellite network case study. Results demonstrate that while PA-CGR improves CGR's congestion avoidance capabilities by leveraging complete path capacity, it still incurs inefficiencies from local traffic bouncing. MG-CGR achieves near-optimal performance with minimal resource usage and delivery delays, rivaling the LP model. In contrast, PCC achieves higher delivery ratios under random traffic but at significant overhead expense. These findings establish MG-CGR as a highly efficient congestion management solution for space-based DTNs, particularly for predictable traffic patterns.

Advanced Congestion Strategies A follow-up journal paper in [MFF18], led by Pablo Madoery, builds upon the concepts from the previous work, with a specific focus on addressing the challenges posed by limited buffer capacities and contact overbook-

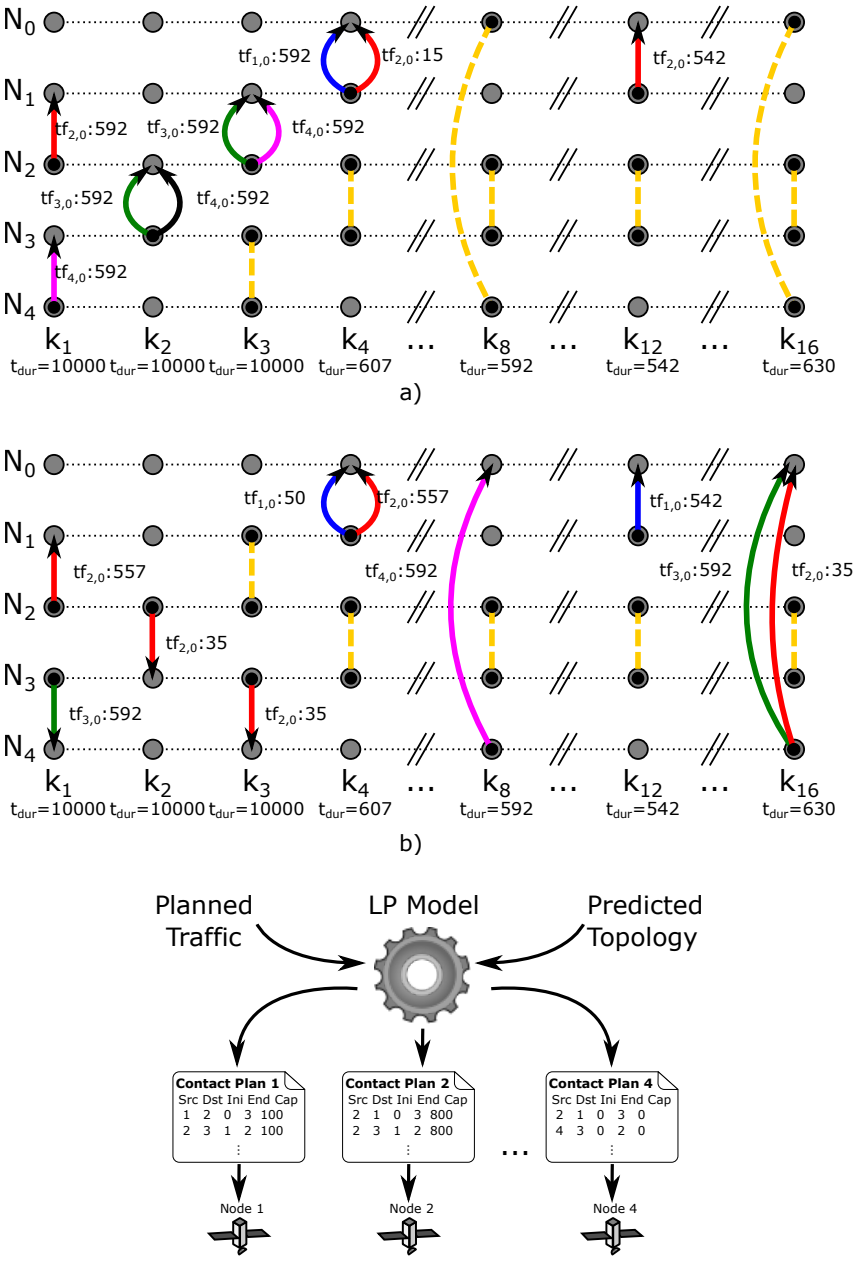

Figure C.18: Congestion management with CGR and MG-CGR. **Top Panel**: Classical CGR allocates flows locally, leading to underutilized contacts (yellow dashed lines) and undelivered traffic (e.g., $tf_{src,dst}$ to $dst = N_0$ from multiple src ranging from node 1 to node 4). **Central Panel**: The LP model optimally forwards traffic using a global view, enabling higher delivery by leveraging all available contacts. **Bottom Panel**: MG-CGR approximates this global optimization by modifying node-specific contact plans, limiting the data volume each node can allocate to ensure efficient use of network resources. (From [Fra+15a])

ing in disruption-tolerant satellite networks. We proposed two novel congestion management strategies: Local Path-Aware Contact Graph Routing (LPA-CGR) and Global Path-Aware Contact Graph Routing (GPA-CGR). LPA-CGR enhances existing CGR by incorporating the complete path capacity in CGR, allowing for more informed local decisions and reducing congestion caused by locally generated traffic. GPA-CGR, on the other hand, leverages global traffic predictability to design custom contact plans for each node proactively. This ensures congestion-free routing while optimizing communication resources. Additionally, the paper introduces an evolutionary optimization approach (EGPA-CGR) to further fine-tune these global strategies based on metrics like delivery ratio and resource usage. Through simulations of a realistic LEO satellite constellation, the study demonstrates that while LPA-CGR outperforms traditional CGR by mitigating local congestion, GPA-CGR and EGPA-CGR achieve superior performance by altogether avoiding congestion, delivering higher payloads with optimized resource utilization. These results highlight the effectiveness of global congestion-aware strategies in managing DTN traffic under predictable conditions.

Learning-Based Techniques A more recent paper in [Yan+24b], led by Lei Yang, tackled the problem of congestion control in deep-space DTNs, focusing on managing extensive propagation delays, energy constraints, and buffer utilization. The proposed solution integrates a delay-sensitive Deep Reinforcement Learning (DRL) approach using a Proximal Policy Optimization (PPO) algorithm. The method models the congestion control problem as a Markov Decision Process (MDP), incorporating delay, buffer utilization, and energy efficiency into the state space. It then uses a reward function to optimize the trade-offs between these metrics. Through simulations involving Earth-Moon and Earth-Mars scenarios, the DRL model demonstrates superior performance compared to baseline methods, achieving minimal packet loss, lower end-to-end delay, and reduced energy consumption. Figure C.19 presents a diagram of the PPO model and the evaluated deep-space scenario. Key findings highlight the model's ability to dynamically adapt transmission rates to traffic patterns. This method mitigates congestion while maintaining energy efficiency and ensuring timely data delivery in challenging deep-space communication environments.

Storage-Congestion Interplay Another critical aspect of congestion management in DTNs is accurately modeling storage capacity and its representation to enable the selection of routes with sufficient capacity for optimal resource utilization. Our in [Ber+24] deals with the problem of modeling and optimizing storage requirements in DTNs, particularly for interplanetary communications, where efficient use of storage and communication resources is critical. In collaboration with Bob Short teams at NASA GRC, we

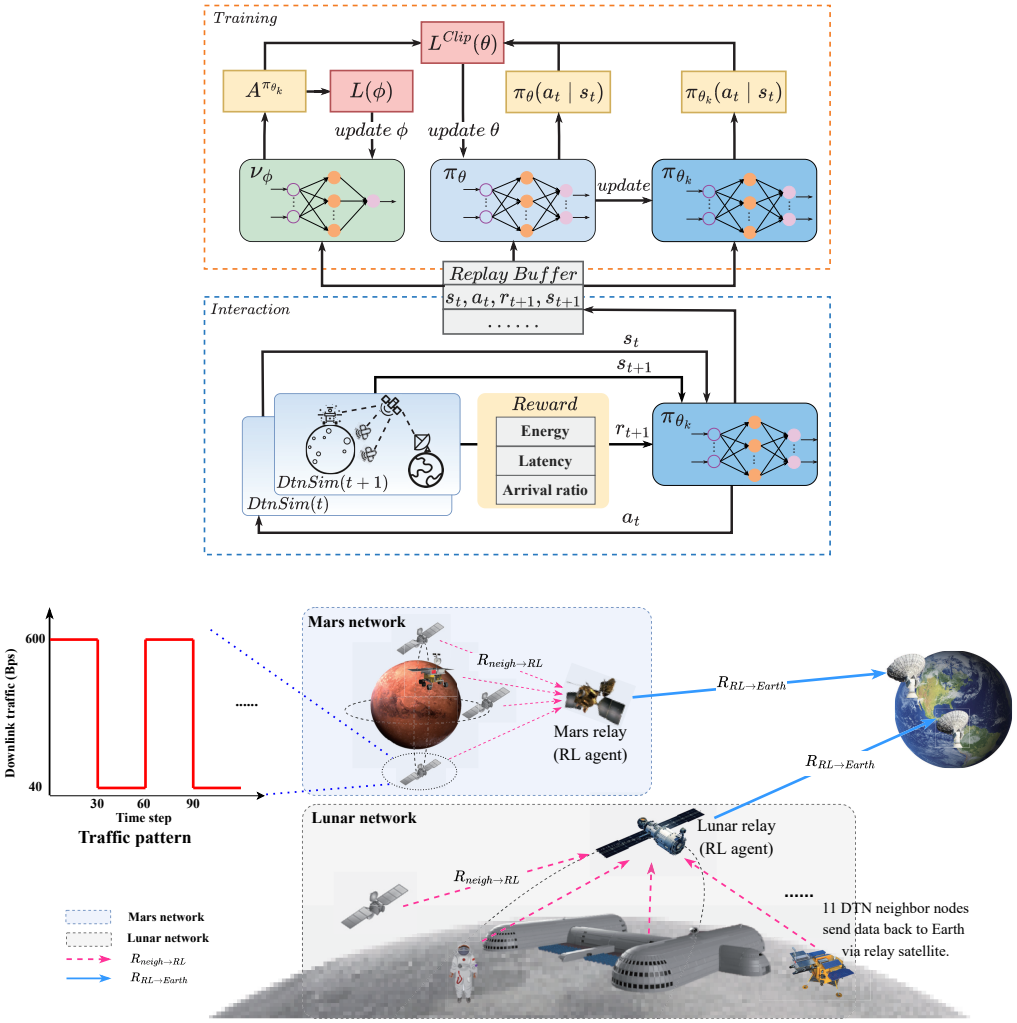

Figure C.19: Proximal Policy Optimization (PPO) model for DTN congestion control. **Top Panel**: The PPO model is structured into two primary phases: interaction and training. During the interaction phase, the agent interacts with the simulated environment (DtnSim) to gather experience data, including states (s_t), actions (a_t), rewards (r_{t+1}), and subsequent states (s_{t+1}), which are stored in a replay buffer. The reward is calculated based on energy consumption, end-to-end latency, and data arrival ratio. The PPO algorithm optimizes two neural networks in the training phase: the policy network (π_θ) and the critic network (ν_ϕ). The policy network refines the strategy for selecting actions, while the critic network evaluates the value of each state-action pair. PPO employs a clipped objective function to stabilize updates and prevent overly aggressive policy changes. The training process iteratively adjusts the policy parameters (θ) and critic parameters (ϕ) using backpropagation until the policy converges to maximize cumulative rewards. **Bottom Panel**: The evaluation scenario involves a deep-space communication system with two primary networks: lunar and Mars. Both networks utilize a relay satellite to transfer data from multiple DTN neighbor nodes to Earth. Data flows are represented by two critical transmission rates: data rate from neighbor nodes to the relay and data rate from the relay to Earth. Traffic patterns alternate between high (600 bps) and low (40 bps) load phases over a 90-step evaluation period to simulate bursts and lulls commonly observed in scientific missions. (From [Yan+24b])

introduced a novel algebraic framework based on semirings to model communication windows, storage capabilities, and transmission durations in DTNs. Using this semiring-based model, they formulate precise methods for computing storage needs to achieve specified transmission durations across a network. Key theoretical contributions include a new semiring structure that encapsulates storage and transmission behaviors and closed-form expressions (referred to as *Nevada constructions*, see Figure C.20) that uniquely define communication paths and their storage requirements using only a few parameters. The paper validates the approach through Earth-Moon and multi-hop interplanetary network simulations, demonstrating how the model can predict storage needs, reduce congestion, and inform network design decisions. The results show that this method significantly simplifies the analysis of storage-constrained DTNs, offering a scalable and computationally efficient approach for real-world applications.

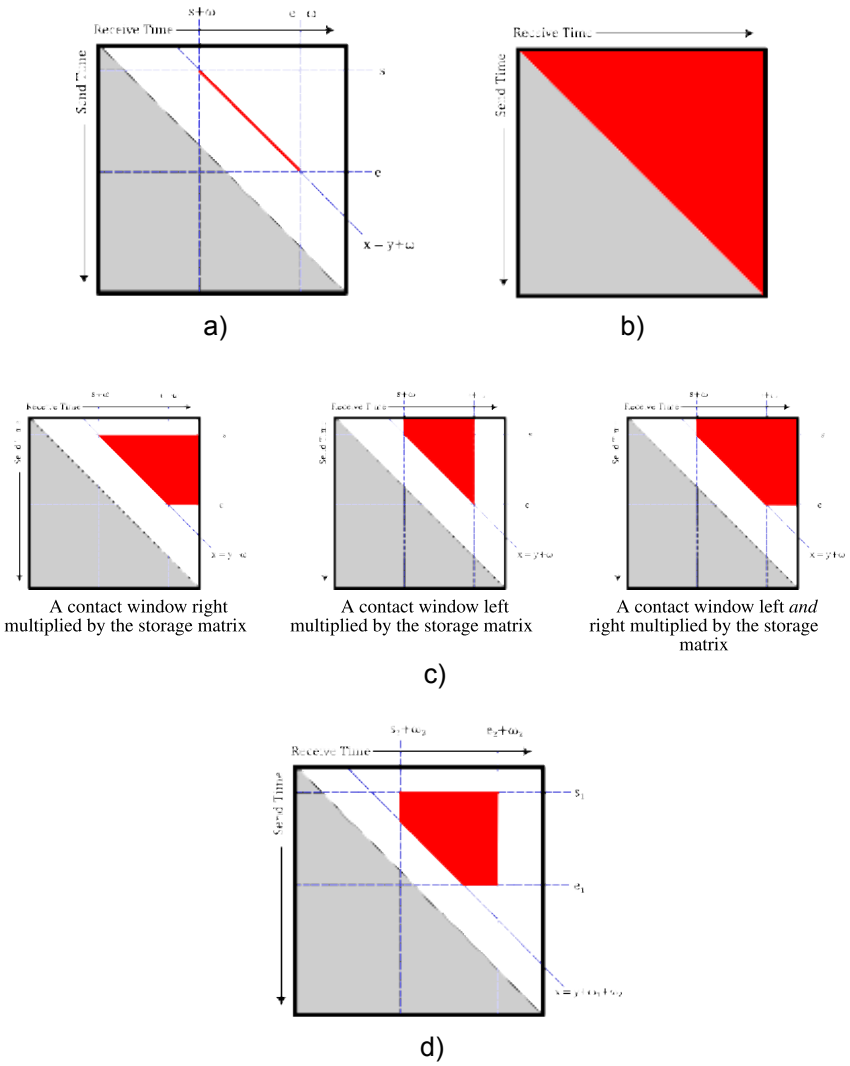

Figure C.20: Nevada constructions. a) **Top Left Panel**: The boolean matrix associated with a contact ($[s, e] : \omega$), where s is the start time, e is the end time of the contact window, and ω is the propagation delay. The red line represents the valid send and receive times within the contact window, constrained by the delay. b) **Top Right Panel**: The storage matrix S, representing the ability to store data at a node. The red region illustrates unlimited storage capability, allowing any sent data to be received at any future time. c) **Center Panel**: Various combinations of contact windows and storage matrices: A contact window right-multiplied by the storage matrix, extending the send-receive times horizontally to accommodate storage (left). A contact window is left-multiplied by the storage matrix, extending the times vertically for received data to be stored (center). A contact window, both left- and right-multiplied by the storage matrix, fully extends the send-receive capabilities' temporal reach (right). d) **Bottom Panel**: The general form of an element of the store-and-forward semiring, $([s_1, e_1] : \omega_1)S([s_2, e_2] : \omega_2)$, visualized as a bounded region (red) defined by five hyperplanes: $y = s_1$, $y = e_1$, $x = s_2 + \omega_2$, $x = e_2 + \omega_2$, and $x = y + \omega_1 + \omega_2$. This representation encapsulates the combined effects of contact windows and storage constraints. (From [Ber+24])

Appendix D

Links and Medium Access Control

This section approximates the question on *How data transfer and multiple access are efficiently coordinated among devices over vast space distances?*

This section shifts focus from multi-hop analyses to the critical challenges of point-to-point communications, diving into contributions at the data link layer, particularly the Medium Access Control (MAC) sublayer. It emphasizes that network layer abstractions discussed earlier are firmly rooted in practical and realistic assumptions of node-to-node data transfers in space systems.

The section is structured into two key sections: **Link** technologies, which explores constraints and valid assumptions while examining how network-layer techniques can enhance link performance, and **Access** considerations, which investigates the overhead imposed by access protocols and how they can benefit from network-layer optimization. Together, these discussions form a cohesive exploration of the foundational mechanisms enabling efficient data transfer and coordination across vast space distances, bridging the gap between theoretical models and real-world implementation challenges.

D.1 Links

This section focuses on my contributions to developing and analyzing communication technologies for point-to-point space links, both for space-to-space and space-to/from-ground communications. Due to the vast distances involved, these links are critical in space networks. Directional transmission methods like high-gain antennas or laser links are required to maximize energy efficiency. My work on link technologies was a foundational stepping stone for the networking research presented in the previous sections, where these links were abstracted into contacts and contact plans. This grounding ensures that the higher-level analyses and protocols central to my research are based on accurate and practical technological assumptions, addressing modern

space communication systems' inherent constraints and capabilities.

My involvement in lower-layer aspects, such as link technologies and physical-layer constraints, was essential to ensuring that higher-level network layer abstractions, like contact plans and routing protocols, are accurate and grounded in real-world feasibility.

This section is organized into two categories: **Radio Frequency** links, detailing the development of the SARE mission ISL modem and its integration with DTN technologies, and **Free-Space Optical** links, which are rapidly emerging as a key enabler for both space-to-ground and space-to-space communications in modern space networks.

D.1.1 Radio Frequency (RF) Links

D.1.1.1 Proximity-1 Modem

My paper in [FK15] addresses the challenge of enabling autonomous and flexible communication for small satellite constellations in LEO, focusing on the limitations of traditional transponders and protocols that lack the adaptability required for inter-satellite links (ISLs) and Earth-satellite links (ESLs). Supported by CONAE personnel, we presented the STI-PRX-01, a modem based on the CCSDS Proximity-1 protocol, to overcome these issues, designed for both ISLs and ESLs. The goal of the modem was to be implemented in CONAE's SARE mission, already discussed in Section A.1.1.1. The modem autonomously negotiates communication parameters such as frequency, modulation, data rate, error correction, and transmission power, adapting dynamically to varying link conditions. It incorporates advanced coding techniques like LDPC, configurable modulation schemes (BPSK, QPSK, and their differential versions), and robust synchronization and equalization methods to handle Doppler shifts and high-frequency offsets. Performance evaluations under stressed scenarios demonstrate significant improvements. With LDPC coding, retransmissions were reduced by 80%, channel usage time decreased, and energy efficiency improved without sacrificing data reliability. The details of the modem, its architecture, and some key performance measurements are provided in Figure D.1. These results highlight the STI-PRX-01's flexibility, autonomy, and suitability for modern small satellite constellations, setting the stage for future plug-and-play communication subsystems.

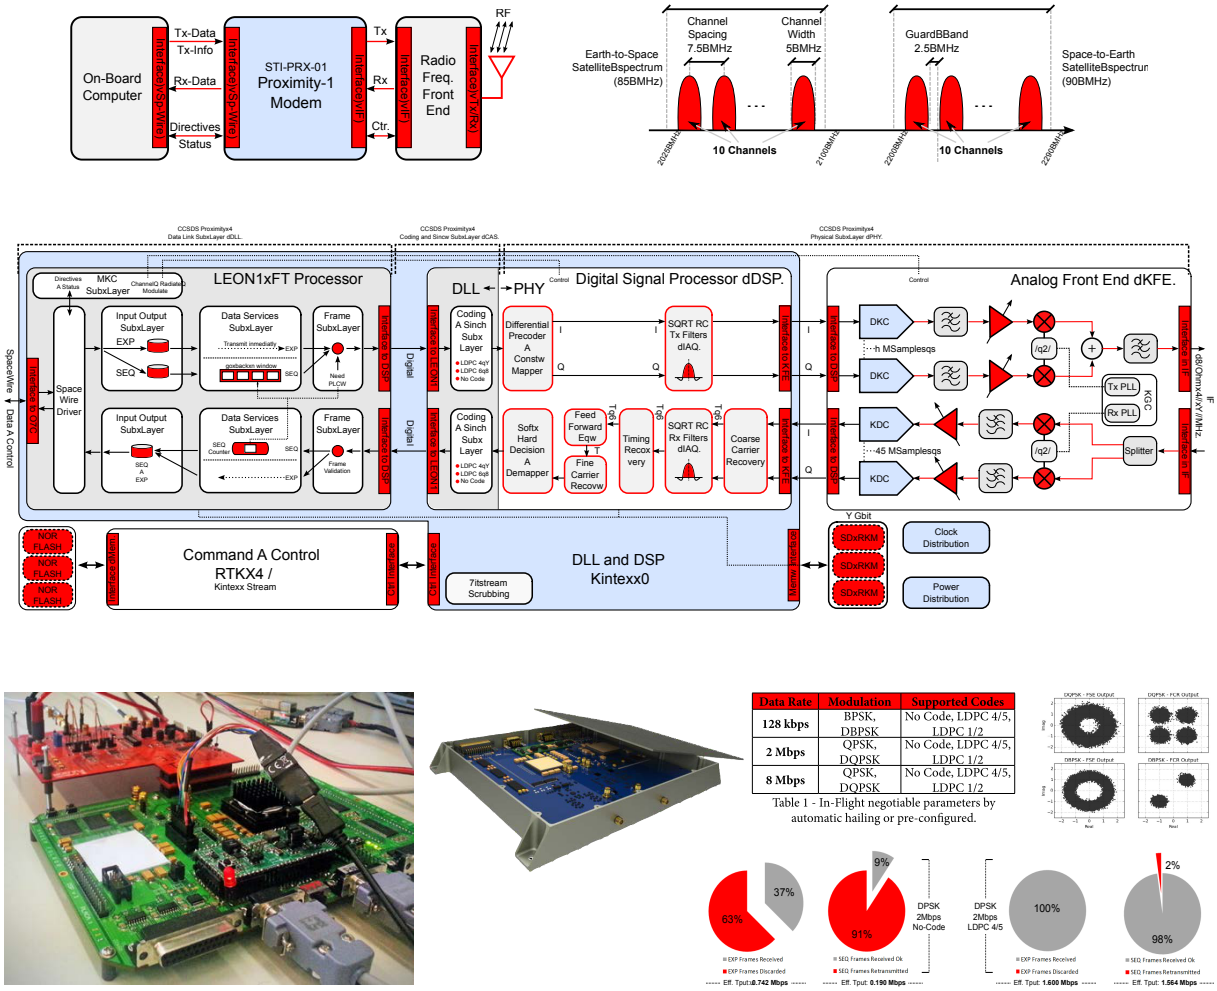

Data Rate	Modulation	Supported Codes
128 kbps	BPSK, DBPSK	No Code, LDPC 4/5, LDPC 1/2
2 Mbps	QPSK, DQPSK	No Code, LDPC 4/5, LDPC 1/2
8 Mbps	QPSK, DQPSK	No Code, LDPC 4/5, LDPC 1/2

Table 1 - In-Flight negotiable parameters by automatic hailing or pre-configured.

Figure D.1: STI-PRX-01 Proximity-1 Modem. **Top Panel**: Context and modem interfaces, illustrating its integration with the onboard computer, radio frequency front-end, and antenna systems. The frequency allocation includes Earth-to-space and space-to-earth bands (658 MHz and 908 MHz) with 10 channels, a channel width of 7.5 MHz, and guard bands of 2.5 MHz for interference-free operation. **Center Panel**: Detailed architecture of the modem, showcasing its functional components: the LEON3-FT processor managing directives and data services, the data-link layer (DLL) performing protocol operations, the digital signal processor (DSP) handling coding, synchronization, and modulation, and the analog front-end (AFE) enabling RF signal processing with coarse and fine carrier recovery, timing recovery, and frequency conversion. **Bottom Panel**: Prototype and finalized product of the STI-PRX-01 modem, accompanied by measurement results. Performance metrics highlight the modem's configurable data rates (128 kbps to 8 Mbps), supported modulations (BPSK, QPSK, and their differential versions), and error correction codes (LDPC). Test results demonstrate high reliability with 98% frame reception success using LDPC 4/5 coding under challenging channel conditions. (From [FK15])

D.1.1.2 Contact Determination and Tracing

This section discusses three works emerging from the COVID-19 pandemic where contact determination and tracing was crucial using RF protocols. These works allowed us to contribute with sophisticated COVID-19 propagation models [Sol+21]. Although not directly related to space, we found that the determination and tracing processes discussed have applicability to the space domain.

Feature selection The paper in [Mad+21] addresses the problem of improving proximity estimation in COVID-19 contact tracing apps that rely on Bluetooth Low Energy (BLE) signals, where existing models fail to accurately determine close contacts due to environmental noise, signal attenuation, and lack of contextual awareness. We proposed a machine learning-based approach to proximity estimation, emphasizing feature selection from BLE signals, such as mean, skewness, and interquartile range. We also incorporated environmental awareness (indoor vs. outdoor) to enhance model accuracy. Experimental campaigns with BLE-enabled smartphones in indoor and outdoor environments evaluated combinations of single and multiple features across logistic regression, support vector machine (SVM), and random forest models. The results demonstrated that models with environmental awareness and multiple features significantly outperformed single-feature benchmarks, achieving proximity detection accuracies of up to 83% indoors and 91% outdoors.

Deep Contact Tracing A second paper in [Che+23] addresses the limitations of traditional contact tracing models that primarily focus on close contacts, which can overlook indirect infection chains in contagious diseases like COVID-19. The proposed method introduces a deep digital contact tracing model that leverages mobile devices to extend contact tracing beyond immediate encounters. The approach uses encounter graphs to identify *deep contacts*, links established through chains of interactions rather than direct proximity. This model incorporates temporal constraints, allowing individuals to compute their contagion risk locally while preserving privacy by minimizing the centralization of contact data. Key outcomes include improved early detection and risk notifications, enabling timely interventions and better monitoring of potential infection paths. The approach's feasibility is demonstrated through simulations and prototype implementations on mobile devices, showing scalable performance while addressing challenges such as data storage, energy efficiency, and privacy concerns. This approach enhances the granularity and effectiveness of digital contact tracing systems for managing epidemics.

D.1.2 Free-Space Optical (FSO) Links

The use of Free-Space Optical (FSO) links in space is not new, but the growing prevalence of networked space missions and the New Space paradigm is driving their accelerated adoption. By concentrating energy into narrow laser beams, FSO links achieve significantly higher data rates per power unit than RF systems. However, their implementation comes with notable challenges, particularly in the precise **pointing requirements** that demand stringent orientation and position control of space platforms. Additionally, for space-to-ground links, **atmospheric conditions**, such as cloud cover, can severely limit the effective use of optical frequencies, posing operational constraints.

> While pointing and atmospheric challenges in FSO can be mitigated at higher networking layers, my focus is on first understanding and optimizing the link layer, enabling the network layer to adapt and maximize the capabilities of the FSO system.

Although not explicitly detailed in the following section, my core experience with FSOs stems from my involvement in the MISSION project in collaboration with Skyloom, a California-based company founded by Argentinian colleagues with ties to the SARE project. Skyloom specializes in developing FSO terminals, and our cooperation focuses on advancing upper-layer technologies to maximize their potential from a network perspective.

The remainder of this section is organized into **Space-to-Ground** and **Space-to-Space** FSO links. In both cases, I leverage collaborations with institutions specializing in the link and physical layers of FSOs to complement my core expertise in the networking layer.

D.1.2.1 Space-to-Ground

FSO Availability This paper in [Gon+20], developed in cooperation with Prof. Kanglian Zhao's team at Nanjing University, addresses the problem of maximizing network availability in FSO communication networks, which are highly susceptible to atmospheric disruptions such as clouds and turbulence. The proposed method introduces a network availability maximization framework based on Integer Linear Programming (ILP) to optimize the placement of FSO relays and the selection of multi-hop paths between nodes. The model considers dynamic weather conditions, link failure probabilities, relay placements, and connectivity constraints. To mitigate the computational complexity of the exhaustive search, a heuristic algorithm is developed that provides near-optimal solutions with significantly reduced runtime. Simulation results demonstrate that the pro-

posed heuristic achieves network availability levels comparable to exhaustive search methods across various network sizes, traffic patterns, and relay deployment scenarios while reducing computational overhead by up to 90%. The study highlights the effectiveness of strategic relay placement and multi-hop routing in maintaining high network availability under adverse atmospheric conditions.

5G FSO Offloading The paper in [Dow+20], developed in cooperation with Alexis Dowhuszko, from Aalto University, addresses the inefficiencies in caching popular content at 5G edge nodes due to the high placing time of terrestrial optical backhaul links, primarily when serving geographically distributed regions with varying content popularity. The proposed method integrates satellite (GEO/LEO) and terrestrial FSO technologies into a hybrid backhaul network to optimize content delivery. The system models GEO satellites in a bent-pipe mode for continuous pan-European coverage and LEO satellites as data mules with sporadic contact opportunities. The approach predicts content requests using correlated Zipf-based popularity distributions and employs a scheduling algorithm to prioritize content delivery during satellite visibility windows. Simulation results show that a hybrid approach significantly reduces placing time compared to terrestrial-only backhauls. Eight-LEO constellations achieve similar performance to GEO at a fraction of the cost. Notably, the placing time for FSO-only backhauls (7.2 hours) is reduced to 5.5 hours with a single LEO, 3.8 hours with a single GEO, and 3.3 hours with 8-LEOs. These findings highlight the potential of small LEO constellations to complement GEO systems and deliver content efficiently and cost-effectively in future 5G networks.

D.1.2.2 Space-to-Space

For the space-to-space FSO application, I began collaborating with i2CAT, based in Barcelona, Spain, in 2023 through a cooperation contract with Inria. Our joint efforts focus on enhancing FSO link performance by optimizing terminal placement and introducing a side channel to accelerate pointing and link establishment processes.

FSO Terminal Placement Our paper in [FRF24], led by myself but in cooperation with i2CAT, addresses the problem of optimizing the placement of FSO terminals on spacecraft to maximize effective contact time for communication links while considering the constraints imposed by gimbal range and motion speed. The study employs two heuristic optimization techniques: Segmented Architecture (SA) and Genetic Algorithms (GA). To explore and optimize placement configurations, these techniques integrate orbital mechanics, spacecraft attitude, and FSO terminal characteristics. Evaluations were

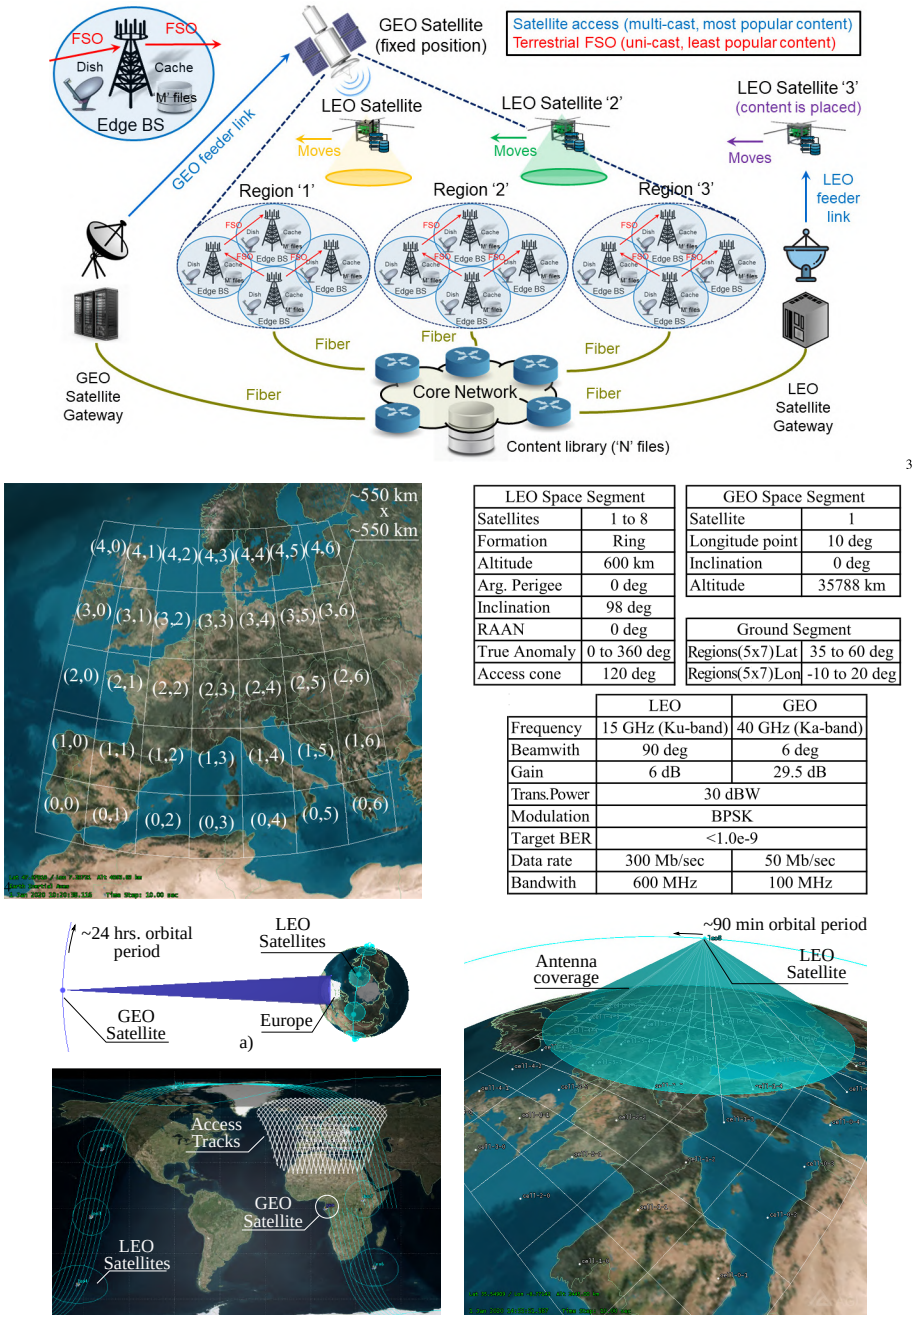

LEO Space Segment	
Satellites	1 to 8
Formation	Ring
Altitude	600 km
Arg. Perigee	0 deg
Inclination	98 deg
RAAN	0 deg
True Anomaly	0 to 360 deg
Access cone	120 deg

GEO Space Segment	
Satellite	1
Longitude point	10 deg
Inclination	0 deg
Altitude	35788 km

Ground Segment	
Regions(5x7)Lat	35 to 60 deg
Regions(5x7)Lon	-10 to 20 deg

	LEO	GEO
Frequency	15 GHz (Ku-band)	40 GHz (Ka-band)
Beamwith	90 deg	6 deg
Gain	6 dB	29.5 dB
Trans.Power	30 dBW	
Modulation	BPSK	
Target BER	<1.0e-9	
Data rate	300 Mb/sec	50 Mb/sec
Bandwith	600 MHz	100 MHz

Figure D.2: Hybrid Backhaul Network for 5G Edge Content Delivery. **Top Panel**: Illustration of the hybrid terrestrial (FSO) and satellite (GEO/LEO) backhaul network for 5G edge caching. The system includes G = 3 geographical regions, L = 3 LEO satellites, and K = 4 edge base stations per region. **Center Panel**: Ground and satellite configuration details, including key parameters such as orbital characteristics, link budgets, and data rates for both LEO (300 Mbps) and GEO (50 Mbps) satellites. **Bottom Panel**: Orbital configurations for the satellite constellations. GEO satellites maintain fixed coverage (blue beam), while LEO satellites provide dynamic coverage (cyan beam). LEO satellites operate in a ring formation, passing over Europe twice daily. (From [Dow+20])

conducted using realistic satellite scenarios to compare baseline configurations with optimized solutions, such as LEO-to-GEO and LEO-to-LEO intersatellite links. The results demonstrate that SA achieves up to a 27.7% improvement in contact time, outperforming GA by a small margin, which also achieved up to 26.8%. Furthermore, both methods ensured robust long-term performance over extended simulation periods, validating their practicality for space communication systems. These findings underscore the significance of using advanced heuristic approaches for optimizing FSO terminal placement in dynamic orbital environments.

Hybrid FSO Evaluation Our second paper in [Fer+24] addresses the challenge of establishing reliable and precise optical inter-satellite links given the narrow field-of-view and high precision requirements of optical terminals, which are often hindered by significant pointing uncertainties in dynamic satellite scenarios. To address this, we proposed a hybrid inter-satellite communication architecture that integrates low-throughput RF and high-throughput optical technologies into a single module. The RF component leverages quasi-omnidirectional antennas to exchange GPS position data in real-time and reduce pointing uncertainties, while the optical component handles high data-rate communication with precise tracking. A hybrid pointing strategy is developed, combining Two Line Element (TLE)-based orbital predictions for coarse pre-alignment with real-time GPS data exchange for refined positioning, followed by optical terminal fine-tuning for link establishment. A custom simulation tool on LEO satellite scenarios evaluates the system's performance. Results demonstrate that the hybrid approach significantly reduces pointing errors. GPS-assisted pointing achieves sub-millidegree precision compared to the multi-degree errors observed in TLE-only strategies after several days of cumulative error. This hybrid strategy effectively minimizes the acquisition time and ensures reliable optical link establishment, addressing critical operational constraints in optical inter-satellite communication.

These papers were developed in collaboration with or led by Elena Fernandez-Niño, a PhD candidate at i2CAT in Barcelona, whom I co-supervise alongside Joan Ruiz-De-Azua. While writing, Elena is advancing a detailed analysis of the gains the RF side-channel provides. She is developing a LoRa-based prototype to validate the concept. The cooperation with i2CAT is highly promising, and I intend to continue and expand this collaboration in the coming years.

Laser Link Scheduler for Deep-Space The paper in [Ger+24b] addresses the challenge of autonomously scheduling free-space laser communications for deep-space networks, particularly for data relay scenarios like those between Mars orbiters, Mars relays, and Earth relays. Manual scheduling approaches fail to scale for modern space

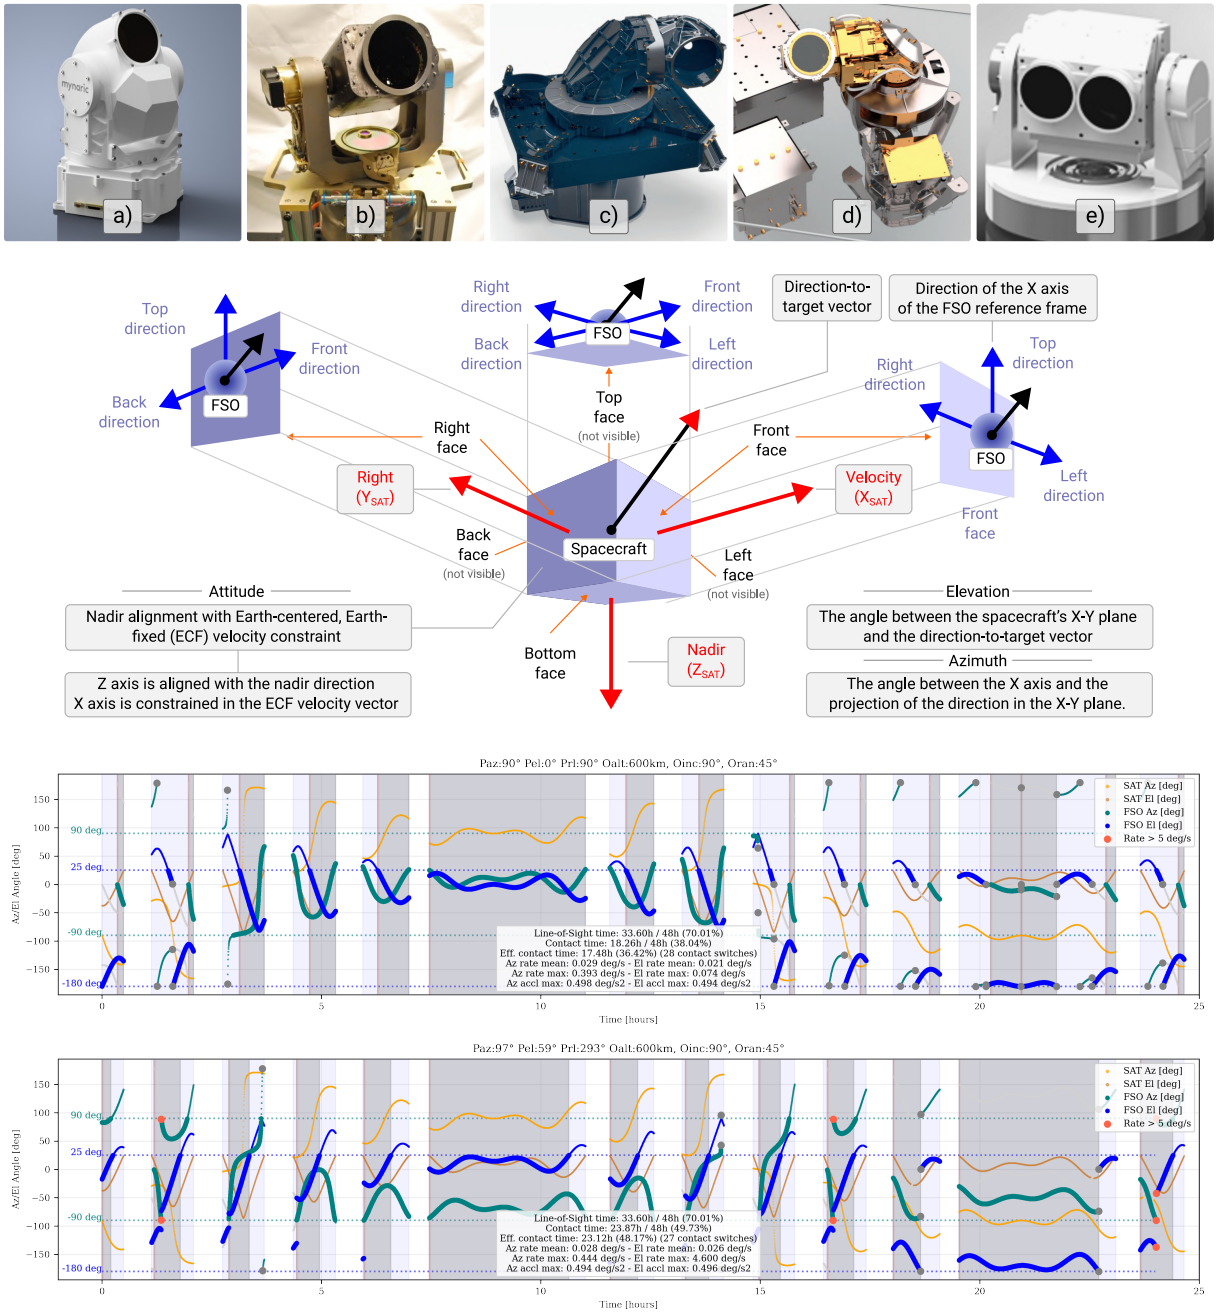

Figure D.3: FSO terminal placement. **Top Panel**: FSO terminals in the market as of 2023: a) Mynaric Mk3 terminal, b) MIT Lincoln Laboratory ILLUMA-T terminal, c) Tesat's LCT 135 GEO terminal (used in EDRS), d) Tesat's SmartLCT 70 LEO terminal, and e) Xenesis Xen-Hub terminal. **Center Panel**: Spacecraft's attitude, reference frame, and baseline placement for FSO terminals. This figure illustrates the default placement of the FSO terminal on the spacecraft, which serves as the baseline solution for comparison with the optimized placements achieved through SA and GA methods. **Bottom Panel**: Timeline plots for best baseline (top) and best SA solutions (bottom) for the 90° inclination orbit. These plots compare the baseline solution, representing the default FSO terminal placement, with the optimized placement achieved through SA. The effective contact time and other performance metrics are shown for both configurations. (From [FRF24])

missions, limiting network capacity and operational flexibility. The proposed Laser Link Scheduler (LLS) leverages a max-flow network capacity model integrated with a novel scheduling algorithm to optimize contact opportunities while ensuring precise Acquisition, Tracking, and Pointing (ATP) for laser communications. The LLS employs a dynamic programming approach to compute edge weights over time-expanded graphs. This enables efficient and synchronized contact scheduling using the Blossom algorithm for maximal weight matching. Key outcomes include a 1.5–2× increase in network capacity compared to existing methods, near-zero buffer wastage, and minimized scheduling delays. These demonstrate LLS's scalability and practicality for large, complex interplanetary networks. The implementation, validated through simulations, highlights the model's ability to optimize resource allocation and enhance the operational efficiency of interplanetary backhaul networks.

D.2 Medium Access Control

Multiple access in space is inherently challenging. As discussed in the previous section, scheduling point-to-point links (e.g., via a contact plan) and concentrating energy efficiently over vast space distances is typically more practical. However, this perspective shifted when I joined Inria and began exploring DtS-IoT, where a different approach to access is required. In DtS-IoT, negotiation and coordination become essential, with some level of scheduling and management by a central entity, but without the rigidity of traditional control channels due to the lack of direct, continuous communication with remote IoT devices. This led me to dive into the fascinating domain of access strategies, which are positioned at the intersection of space predictability and the reduced-control, decentralized nature of IoT negotiations.

> While **multiple access** over vast space distances is generally impractical, the direct-to-satellite access domain necessitates the development of minimal negotiation mechanisms to share a medium spanning extensive coverage areas efficiently.

This section summarizes my contributions to **Direct-to-Satellite IoT**, encompassing analyses, enhancements, and optimizations of existing LPWAN protocols. It then explores NTN management and operations, the integration of IoT access with DTN, and concludes with a discussion on multiple access, extending beyond the IoT domain into the cellular domain.

D.2.1 Direct-to-Satellite IoT Access

This section highlights one of my career's most extensively developed areas: IoT access in DtS-IoT systems. My approach centers on leveraging standardized LPWAN protocols, such as LoRa/LoRaWAN and NB-IoT, and adapting them to the unique dynamics of LEO satellite operations. This strategy capitalizes on Commercial of the Shelf (COTS) devices and chipsets, benefiting from the cost efficiencies of mass-market production. In satellite IoT missions, the affordability of ground devices is a critical factor in the overall system cost equation. This section comprehensively explores DtS-IoT access, including detailed analyses, protocol enhancements, access policies, and management strategies. Additionally, it discusses the development of custom protocols tailored to standard LoRa modulations, offering novel solutions to address the challenges of satellite-based IoT connectivity.

D.2.1.1 Analysis

My most comprehensive analysis of LPWAN is presented in [FIV22], discussed in Section A.2.1.2, where an LPWAN-based architecture for DtS-IoT was proposed. Also, the initial survey in [FCA19] was discussed in that section. This section complements that work by including two additional papers that delve deeper into analyzing and evaluating LPWAN technologies.

NewSpace LPWAN survey The paper in [LLF24], developed in cooperation with Oscar Ledesma and Paula Lamo, analyzes the integration of LPWAN into LEO satellite constellations within the evolving NewSpace paradigm. The problem addressed revolves around achieving scalable, energy-efficient, and cost-effective global IoT connectivity, especially in remote regions, while overcoming the challenges of spectrum efficiency, interference, and limited bandwidth associated with LPWAN technologies in space. We analyze the potential of LPWAN protocols such as LoRa/LoRaWAN and NB-IoT for satellite communications, comparing their operational characteristics, including modulation schemes, data rates, and energy efficiency. The study highlights key innovations like Long-Range Frequency-Hopping Spread Spectrum (LR-FHSS) for enhancing scalability and collision resistance and examines adaptations of NB-IoT under 3GPP standards for satellite compatibility. Simulation results underscore the suitability of LoRa/LoRaWAN for low-power IoT applications and the advanced capabilities of NB-IoT in high-bandwidth scenarios. The work provides a comprehensive framework for designing hybrid terrestrial-satellite IoT systems. It addresses coverage, scalability, and security limitations while identifying future trends, such as integrating AI and blockchain for autonomous resource management and secure communications.

LoRaWAN Performance Analysis Our paper in [Sca+24], led by Germán Scapin, a Master's student in the context of the MISSION project, addresses the challenges of optimizing DtS-IoT networks using LoRaWAN technology, focusing on key aspects such as operation modes and network server placement. To tackle the unique constraints of DtS-IoT, including intermittent connectivity and long transmission distances, the study evaluates various LoRa operation modes characterized by beacon intervals and spreading factor configurations, assessing their impact on data extraction rate (DER) and goodput. Additionally, it investigates the placement of the network server—either on the satellite or on the ground—considering latency, processing overhead, and resource utilization. Using simulations with tools like LoRa-Space and FLoRaSat, the study reveals that increasing the number of communication channels significantly improves DER and goodput by reducing collisions. Variable spreading factors provide performance gains but require fine-tuned management for optimal results. On server placement, onboard network servers enable faster acknowledgments and reduced latency but demand more significant satellite resources. In contrast, ground-based servers offer better scalability and flexibility at the cost of higher delay. The results underscore the importance of tailored operation modes and strategic server placement in maximizing the performance of LoRaWAN DtS-IoT networks.

Machine Learning (ML) Our recent paper in [Cho+24], led by Alexander Ylnner, addresses the challenge of integrating Machine Learning (ML) into DtS-IoT and Satellite IoT systems in general to optimize various aspects such as resource allocation, routing, and security, which are crucial for scalable and efficient satellite-IoT operations. We introduced a taxonomy that categorizes ML techniques like Deep Learning, Reinforcement Learning, and Federated Learning according to their applications in SIoT, emphasizing their roles in managing radio access, network resources, and application-specific tasks (see Figure D.4). Key findings reveal critical gaps, including the lack of differentiation between on-ground and space-based processing, limited incorporation of SIoT-specific metrics, and underestimation of satellite resource constraints. The paper concludes by identifying research directions, such as leveraging advanced ML methods (e.g., Graph Neural Networks and Federated Learning), addressing privacy and security, and enhancing space-terrestrial integration for next-generation networks. It provides a roadmap for aligning ML research with industrial SIoT needs.

D.2.1.2 Enhancement

This section dives into the specifics of LR-FHSS. Developed by Semtech, LR-FHSS is an advanced modulation scheme designed to enhance the scalability and robustness of legacy LoRa technology. As depicted in Figure D.5, it splits data transmissions across

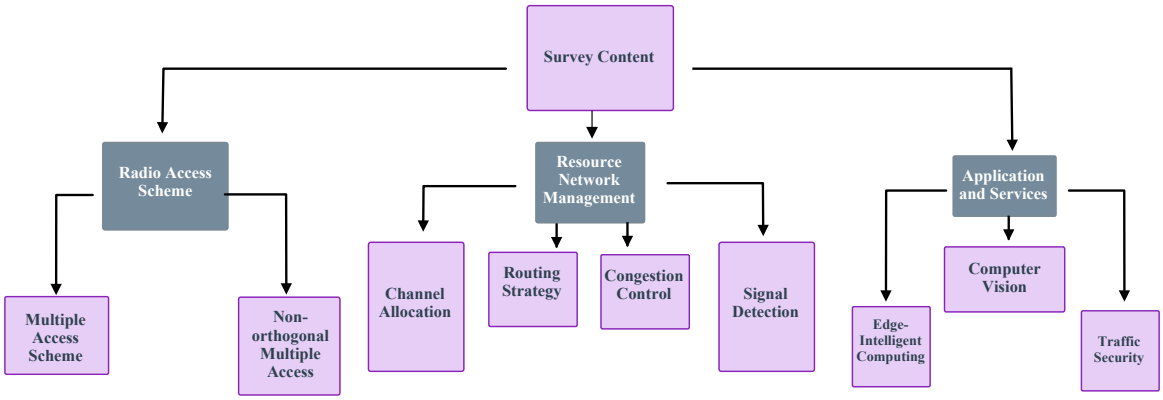

Figure D.4: Taxonomy of survey content for Satellite Internet of Things SIoT, categorizing key areas into Radio Access Schemes, Resource Network Management, and Applications and Services. Subcategories highlight critical components such as multiple access schemes, channel allocation, congestion control, signal detection, edge-intelligent computing, and traffic security, outlining the scope of ML applications in SIoT systems. (From [Cho+24])

multiple frequencies and time slots, reducing collisions and improving spectral efficiency. This makes LR-FHSS particularly well-suited for applications requiring high device density and long-range communication, such as satellite-based IoT networks.

Satellite IoT has been my primary integration project at Inria. In collaboration with Oana Iova, Fabrice Valois, and Prof. Hervé Rivano from Inria Agora, we combined their LPWAN expertise with my satellite knowledge to advance LR-FHSS. I supervise Diego Maldonado, who leads this effort, focusing on optimizing the protocol for improved performance in DtS-IoT applications. This work cooperates with Alexandre Guitton (Université Clermont Auvergne, France) and Megumi Kaneko (National Institute of Informatics, Japan).

The remainder of this section considers our contribution for **headerless frame recovery** and **sequence design and demodulator allocation**, both for LR-FHSS

Headerless Frame Recovery The headerless frame recovery research progressed in two stages. In the first stage, I led the effort, supported by Oana Iova and Alexandre Guitton, to assess the feasibility and validity of the payload selection method. Once its applicability was confirmed, the research was delegated to Diego Maldonado, who received additional guidance from Megumi Kaneko during his internship at National Institute of Informatics.

My first paper in [FGI23a], initially previewed in [FGI23b] and led by me as an early

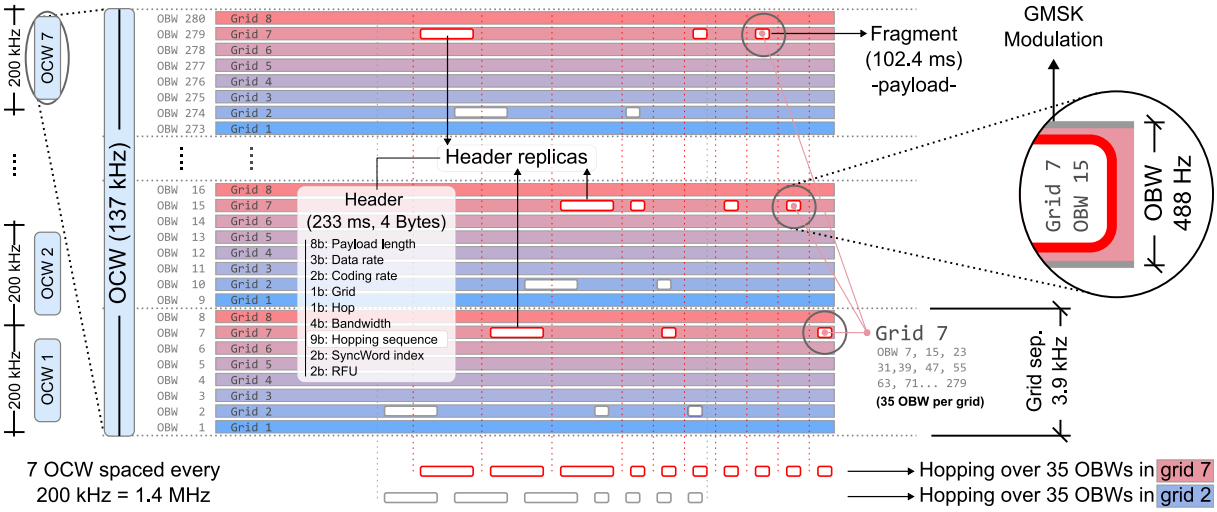

Figure D.5: LR-FHSS Time-Frequency Representation. The figure illustrates the LR-FHSS frame structure, showing frequency hopping across Occupied Channel Widths (OCWs) and grids. Each grid contains 35 Orthogonal Bandwidths (OBWs) spaced 3.9 kHz apart, with 7 OCWs distributed across 1.4 MHz bandwidth. Headers (233 ms, 4 bytes) and payload fragments (102.4 ms) perform frequency hopping within their assigned grid to maximize reliability. Header replicas improve decoding probability but increase congestion, leading to frame loss under high traffic, particularly due to collisions in overlapping OBWs. The inset highlights the GMSK-modulated OBWs within a single grid (e.g., Grid 7). (From [FGI23a])

exploration of the topic, tackles the reliability challenges of Long-Range Frequency-Hopping Spread Spectrum (LR-FHSS) communications in satellite IoT networks, specifically addressing header loss—a critical factor contributing to frame loss under high traffic conditions. To tackle this, the study proposes a novel headerless frame recovery methodology. The approach includes (a) a mathematical model to analyze the fragility of LR-FHSS headers, (b) an MILP model to identify and extract headerless frames, and (c) a computationally efficient heuristic algorithm designed for resource-constrained processors on LEO satellites. Extensive simulations demonstrate that both the MILP model and heuristic achieve 100% true positive rates in detecting headerless frames, with the heuristic offering optimal performance at significantly reduced computational costs. Furthermore, the proposed methods deliver substantial improvements in frame extraction probabilities compared to traditional LR-FHSS, making the techniques highly suitable for managing high traffic loads in resource-limited satellite IoT systems. Key results show that this headerless recovery approach enhances throughput, mitigates the impact of collisions, and optimizes system resource utilization.

Our second contribution in [Mal+24a] to this topic was led by Diego Maldonado, my PhD student at Inria. This paper extends the previous work on LR-FHSS for DtS-IoT networks by adding more realistic considerations to the headerless decoding strategy and

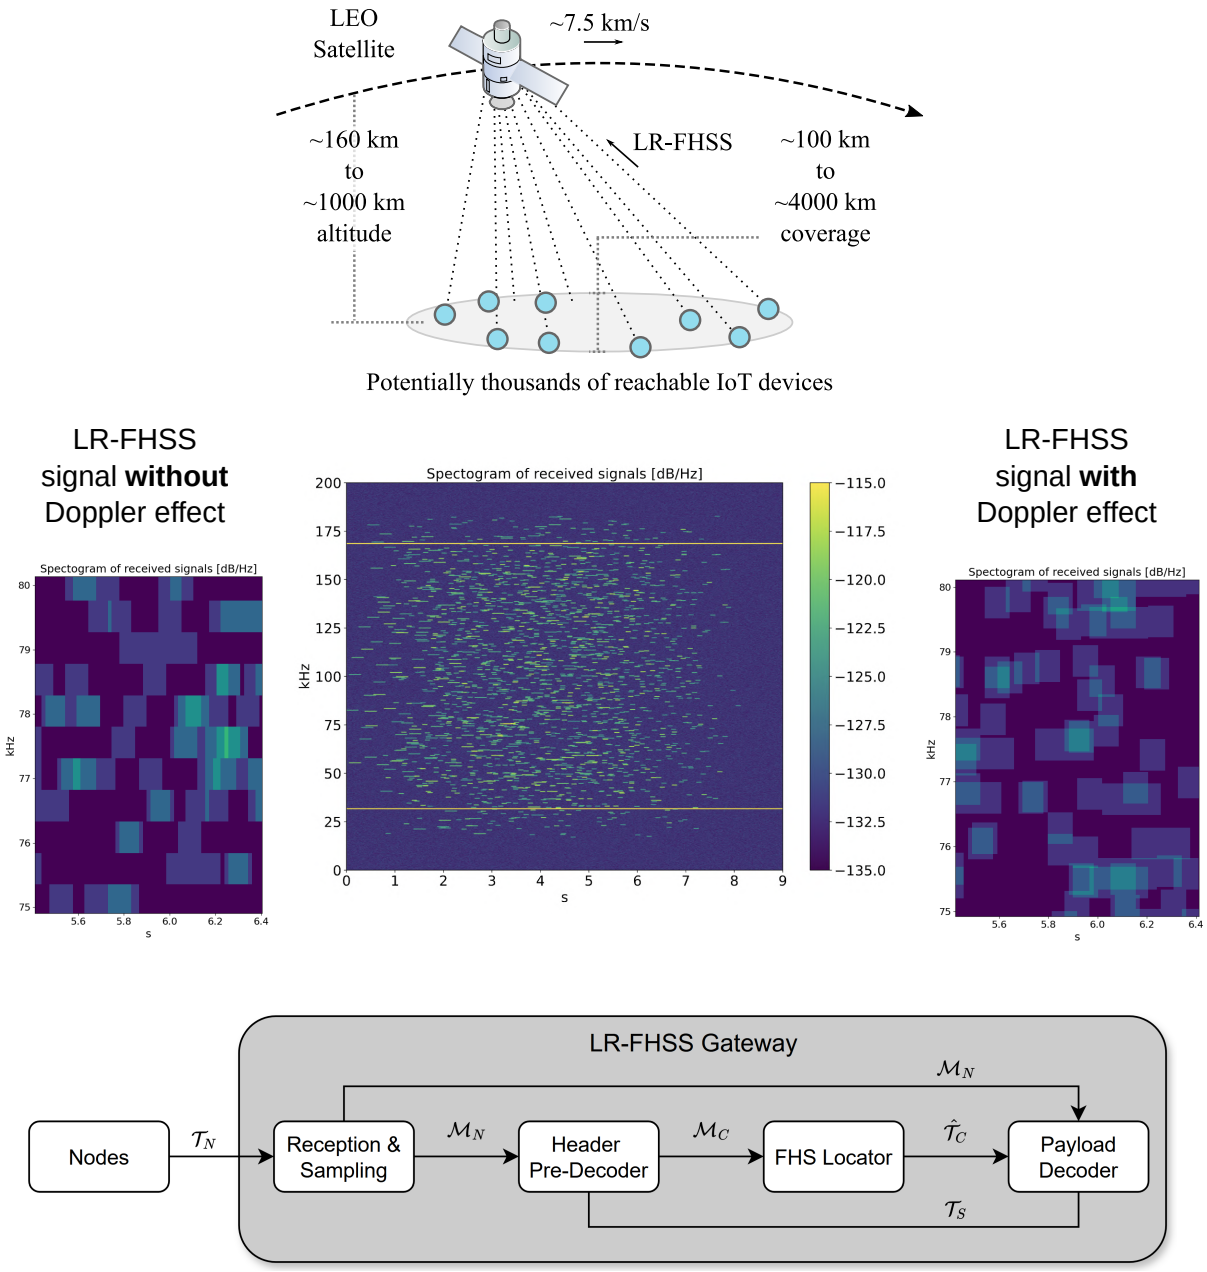

Figure D.6: LR-FHSS Decoding in the DtS-IoT Context. **Top Panel**: Illustration of LR-FHSS deployment in the DtS-IoT context, showing a LEO satellite traveling at 7.5 km/s with coverage extending from 100 km to 4000 km, enabling communication with thousands of IoT devices. **Center Panel**: Spectrogram comparison of 100 received LR-FHSS transmissions with (right) and without (left) Doppler effects. The Doppler effect significantly distorts the frequency hopping patterns, complicating the selection of payload segments and increasing the challenge of frame decoding. **Bottom Panel**: LR-FHSS gateway receiver architecture incorporating headerless decoding capabilities. The design features modules for reception and sampling (T_N), header pre-decoding (M_N), frequency hopping sequence (FHS) locator (M_C), and payload decoding (T_S), optimized for Doppler-compensated processing in satellite-based scenarios. (From [Mal+24a])

framing it into a realistic LR-FHSS receiver architecture that incorporates Doppler compensation techniques. The receiver model is equipped with modules for precise header decoding, a Doppler-aware frequency hopping sequence (FHS) locator, and an exhaustive payload decoding process that accounts for realistic satellite-terrestrial communication conditions, such as variable payload sizes and partial header losses. A diagram of the receiver, as well as context about LR-FHSS in a Doppler-sensitive scenario, is presented in Figure D.6. Simulation results generalize my early work and demonstrate that the enhanced receiver significantly improves the frame decoding rate, achieving up to a 50% increase in terrestrial and satellite scenarios. The system performs well under varying levels of network traffic, with a near-perfect F1 score in satellite scenarios for low to moderate transmission densities. This underlines its scalability and effectiveness in mitigating collision-induced data loss. These advancements pave the way for more efficient and reliable DtS-IoT communications, enhancing global connectivity for low-power IoT devices.

Sequence Design and Demodulator Allocation The paper in [Mal+24b], currently under review but available online, also addresses the scalability challenges of LR-FHSS modulation for DtS-IoT networks, but from a different perspective. It supports massive device connectivity and efficiently utilizes Frequency Hopping Sequences (FHS) and demodulator resources. To tackle these issues, Diego Maldonado proposed two main contributions: 1) the introduction of optimized Wide-Gap FHS (WGFHS) tailored for LR-FHSS modulation, which ensures better Hamming correlation properties and reduces interference among overlapping transmissions, and 2) the development of two novel demodulator resource management strategies—Early-Drop, which releases demodulators when decoding is deemed unachievable, and Early-Decode, which initiates decoding once the minimum fragments required for successful recovery are available. Figure D.7 illustrates the early drop and early decode strategies and the impact of the combined FHS sequence in the collision rate of LR-FHSS. Extensive simulations validate the approach, demonstrating up to 50% improvement in decoding capacity under specific configurations. The study highlights the synergy between advanced FHS design and intelligent resource management. Pushing the scalability limits of LR-FHSS will lead to robust and efficient DtS-IoT communication networks.

D.2.1.3 Transmission Policies

The paper in [Álv+22], led by Guido Alvarez, from CONAE, addresses the scalability and energy efficiency challenges of uplink transmission in LoRa-based DtS-IoT systems, including LoRaWAN and LR-FHSS (called LoRa-E back in the days). These challenges are exacerbated by the dynamic nature of LEO satellite passes and the limited capacity

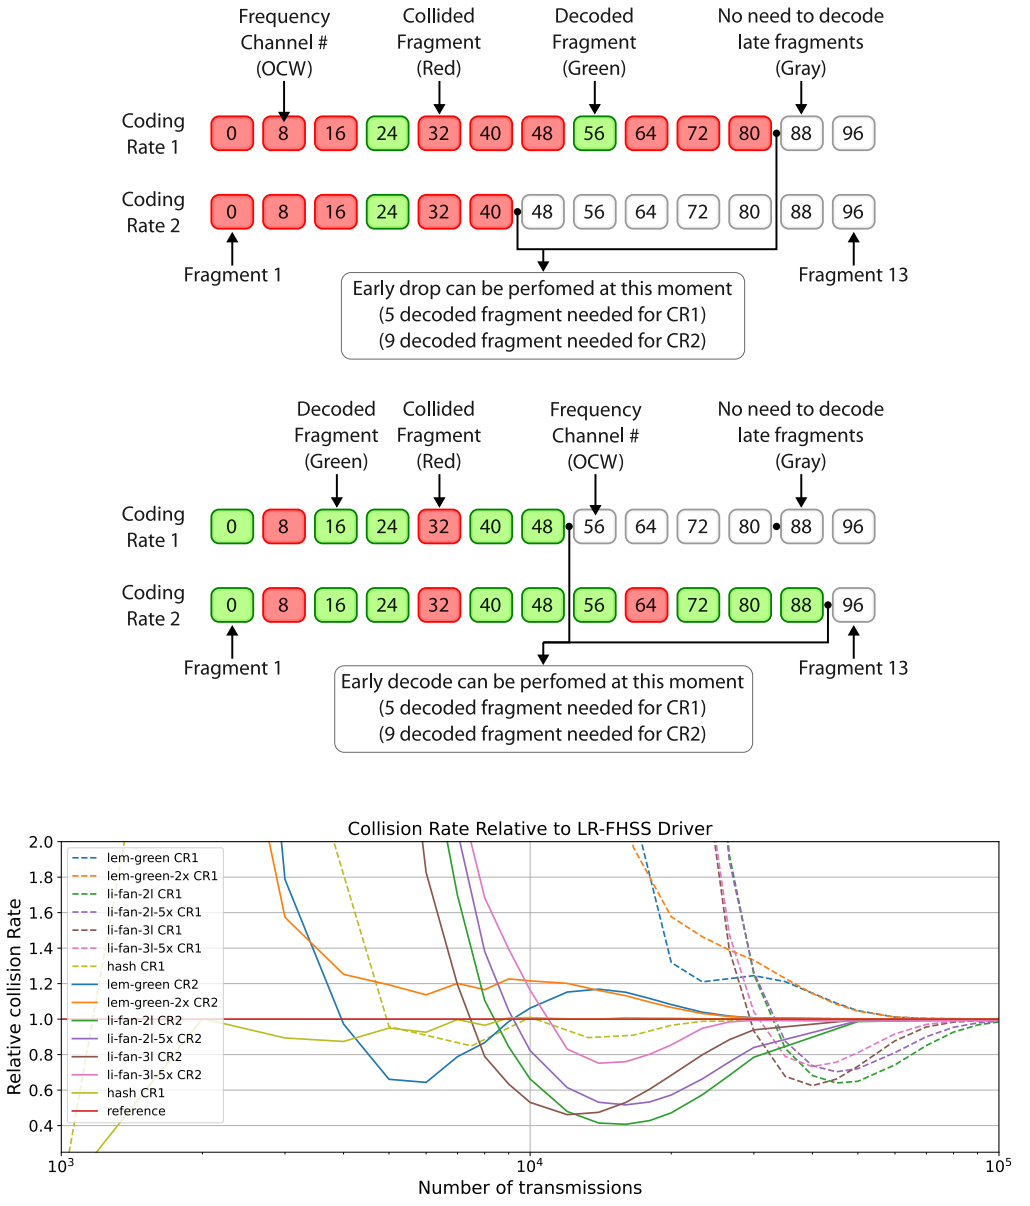

Figure D.7: LR-FHSS demodulator allocation and sequence design. **Top Panel**: Illustration of the Early Drop strategy for a 13-fragment payload. Decoded fragments (green) and collided fragments (red) are tracked across frequency channels (OCWs). The Early Drop mechanism releases demodulation resources once the minimum number of fragments required for decoding (5 for CR1 and 9 for CR2) is reached, avoiding unnecessary decoding of later fragments (gray). **Center Panel**: Representation of the Early Decode strategy for the same 13-fragment payload. The decoding process begins when the minimum fragments necessary for successful recovery (5 for CR1 and 9 for CR2) are collected, optimizing resource utilization and reducing latency. **Bottom Panel**: Relative collision rate for different LR-FHSS FHS families under configurations without granularity, no headers, and unlimited decoding capacity. Strategies include: LEM-GREEN: Wide-gap sequences with collision-aware optimization. LEM-GREEN-2X: Double-density wide-gap sequences. LI-FAN-2X: Optimized sequences with reduced overlap in frequency allocations. LI-FAN-3X: High-density sequences for improved scalability. HASH: Randomized sequences for uniform collision distribution. (From [Mal+24b])

of traditional LoRaWAN Aloha-based MAC protocols, which suffer from low scalability and high collision rates in dense deployments. The paper proposes novel uplink transmission policies that leverage satellite trajectory information to optimize channel access. These policies include trajectory-based strategies, in which devices dynamically select transmission parameters, such as spreading factors (SFs) and data rates (DRs), based on satellite position and channel conditions. Additionally, the paper introduces skip-based schemes, which allow devices to postpone transmissions during congested periods to take advantage of better future opportunities. An optimal MILP model is also formulated to establish an upper bound on performance by computing theoretical collision-free and energy-efficient schedules. The policies are evaluated using extensive simulations that incorporate realistic satellite pass dynamics. Results demonstrate that trajectory-based policies significantly improve network scalability, doubling the number of supported devices compared to baseline approaches. Skip-based policies further enhance scalability by up to 30% under high traffic loads. The MILP model confirms the potential for optimal scheduling to maximize channel utilization while minimizing energy consumption. Additionally, LR-FHSS modulation increases network scalability by a factor of 75 compared to legacy LoRa, albeit with a 30% increase in power consumption per device. These findings highlight the feasibility of scalable and energy-efficient DtS-IoT systems, with trajectory-based approaches effectively balancing performance trade-offs.

The paper in [VFH21], led by Kai Vogelgesang, a Master's student I supervised at Saarland University, introduces a novel Transmission Probability Function (TPF) to optimize uplink communication for LoRa-based DtS-IoT systems under the dynamics of LEO satellites. This function perfectly fits the transmission policy architecture presented in the paper discussed above. The problem addressed revolves around managing the MAC layer in DtS-IoT systems, where thousands of IoT devices simultaneously contend for limited uplink opportunities during a satellite pass, resulting in frequent collisions and wasted transmissions. The proposed method utilizes a framed slotted Aloha model, compatible with LoRaWAN, to dynamically control the number of contending devices. The TPF calculates the optimal transmission probability for devices based on the estimated cardinality of the collision set and beacon information, enabling efficient slot utilization. Additionally, the study addresses a hidden-terminal problem caused by satellite dynamics, introducing two solutions: estimator throttling, which adjusts the collision set estimate, and perceptive devices, which use trajectory data to avoid wasted transmissions. Simulations in realistic satellite scenarios show that TPF achieves up to 87% of theoretical throughput, significantly outperforming baseline methods. Figure D.8 depicts the scenario and explains the results across different satellite pass types. The proposed approaches effectively mitigate contention issues, improve extraction rates, and reduce energy loss, marking a substantial step forward in

scalable DtS-IoT communication.

D.2.1.4 Custom Protocol Design

This section presents the design of custom link layer protocols and related analysis frameworks based on existing LoRa physical layers.

RESS-IoT Our paper in [Ort+21], led by Raydel Ortigueira, a PhD student I co-supervised with Prof. Sandra Céspedes, addresses the challenge of designing an energy-efficient and scalable MAC protocol for DtS-IoT systems, where low-cost nanosatellites (CubeSats) connect to ground sensor nodes constrained by limited power and processing capabilities. The proposed Reserve-and-Send Sift-IoT (RESS-IoT) protocol integrates the LoRa physical layer with a novel MAC layer mechanism inspired by the SIFT protocol, employing a two-phase communication round: a reservation phase with variable-length Clear-to-Send (VCTS) frames to allocate transmission slots, and a transmission phase to optimize data delivery. Figure D.9 depicts the phases graphically. The analytical model and simulations reveal significant performance improvements. Compared to state-of-the-art protocols such as Enhanced Aloha (EA) and Unconfirmed Class A LoRaWAN (UCAL), RESS-IoT significantly reduces energy consumption by up to 409% in satellites and up to 716% in nodes for large-scale networks, while achieving up to 926% higher throughput in 1000-node scenarios. Additionally, the protocol maintains high fairness in channel access and efficiently manages collisions in dense networks. These results demonstrate the potential of RESS-IoT to support scalable and energy-efficient DtS-IoT deployments.

MSDQ Our paper in [GFC24], led by Jason Gerard, a former Master's student and now a PhD candidate co-supervised by myself and Prof. Sandra Céspedes from Concordia University in Montreal, Canada, tackles the inefficiencies of existing MAC protocols in supporting large-scale DtS-IoT networks. The work addresses critical energy efficiency, scalability, and protocol simplicity challenges in these resource-constrained environments. We propose the Massively Scalable Distributed Queuing (MSDQ) protocol, which builds upon Distributed Queuing (DQ) and RESS-IoT concepts to address these limitations. MSDQ introduces key improvements, including a reverse-exponential distribution for slot reservation to reduce collisions, a contention tree algorithm (CTA) for efficient collision resolution, and logical queues for streamlined data transmission and contention handling. Extensive simulations, summarized in Figure D.10, using the FLoRaSat simulator demonstrate that MSDQ outperforms Distributed Queuing (DQ) and RESS-IoT in critical performance metrics. Specifically, MSDQ achieves over 2.2 times

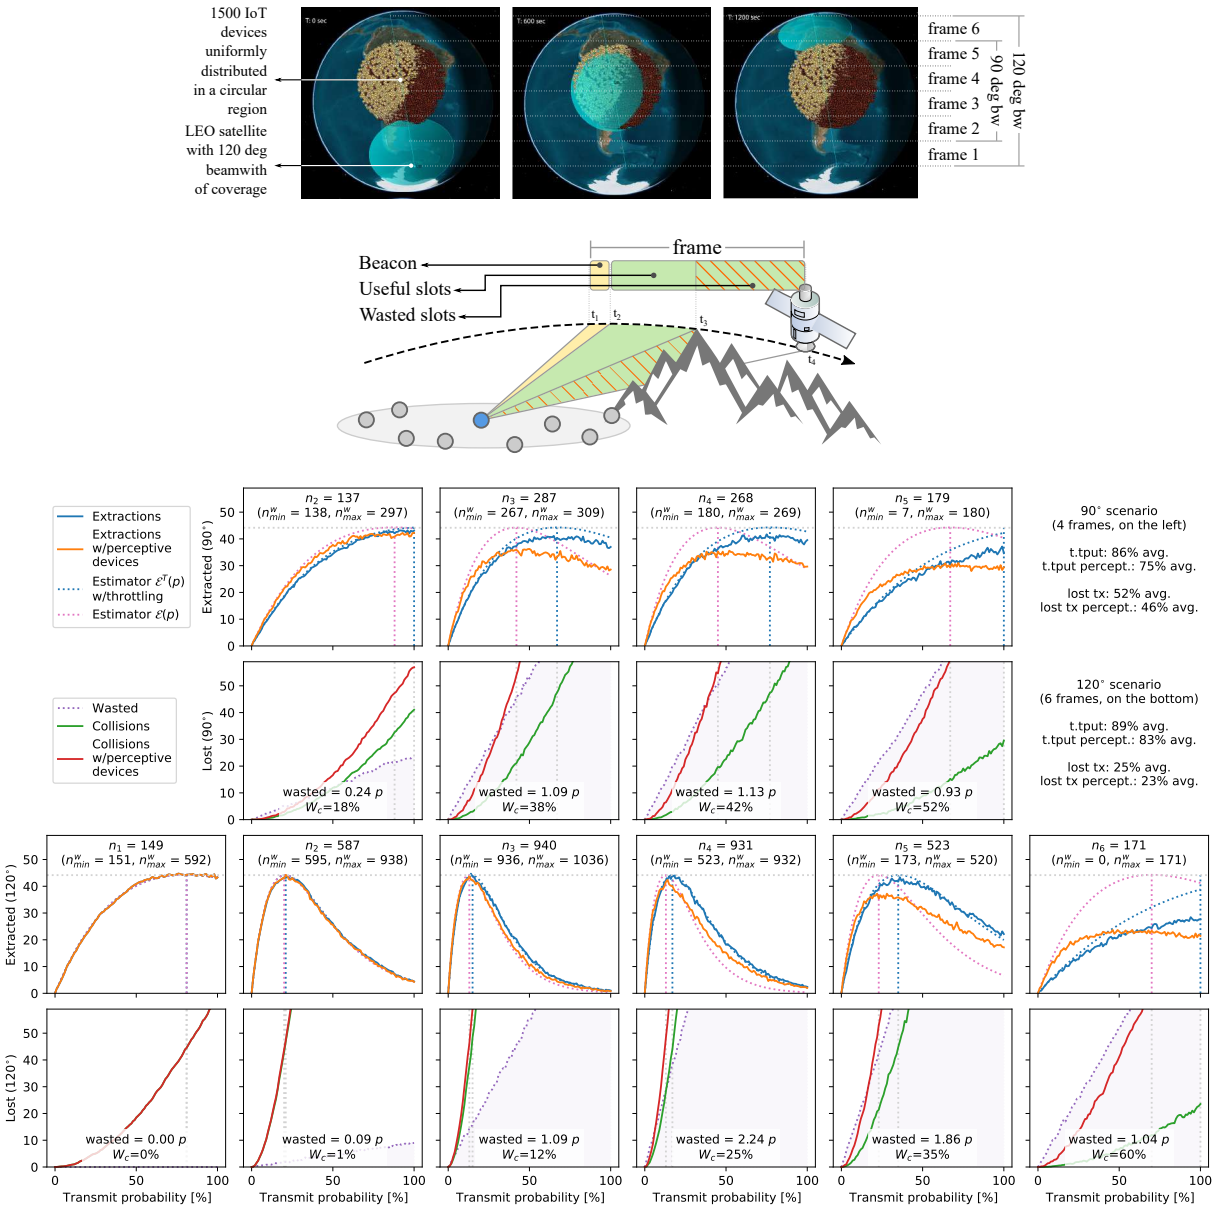

Figure D.8: Transmission Probability Function (TPF). **Top Panel**: Case study scenario of DtS-IoT with a LEO satellite covering IoT devices using 90° and 120° beamwidths. **Center Panel**: Illustration of useful and wasted slots within a frame. The elevation profile (depicted as mountains) determines the visibility horizon for a specific device (blue). Devices beyond the visible horizon during certain slots contribute to wasted resources. **Bottom Panel**: Performance metrics for 90° (top row) and 120° (bottom row) beamwidth scenarios, with extraction rates, collision metrics, and lost messages. Key insights include: *Vertical Dotted Lines*: Optimal transmission probabilities. *Horizontal Dotted Lines*: Theoretical throughput upper bound for slotted Aloha with 120 slots (36.7%, or 44 extracted packets). *Reachable Devices*: n_f indicates devices reachable in the beacon, while $n_{\min}$ and $n_{\max}$ represent the minimum and maximum reachable devices across slots. The averaged throughput and percentage of lost transmissions are summarized in the top-right corner. (From [VFH21])

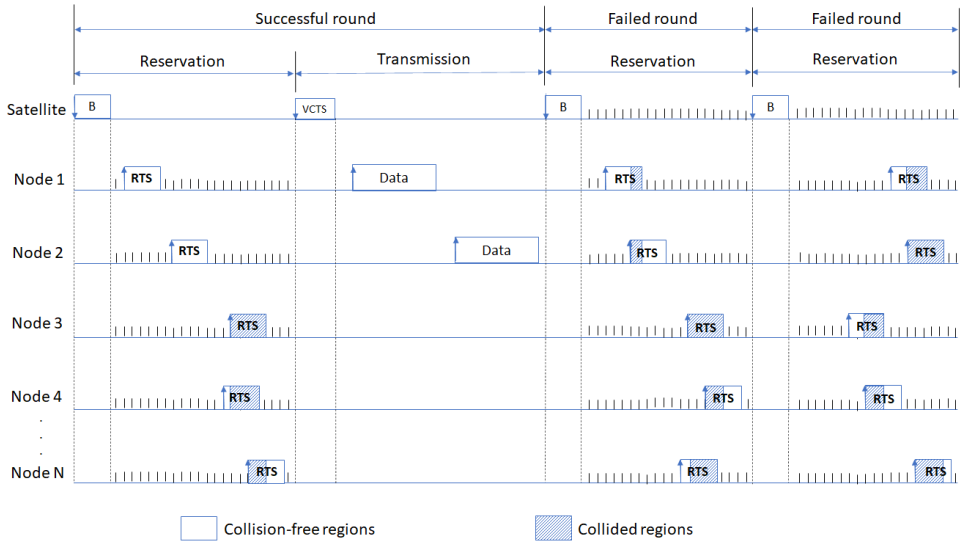

Figure D.9: RESS-IoT Protocol Overview. Two-phase operation of the RESS-IoT protocol: Reservation Phase and Transmission Phase. **Reservation Phase**: IoT nodes send Request to Send (RTS) messages to the satellite during designated slots. Collision-free RTS messages (white regions) allow nodes to reserve transmission slots, while collided RTS messages (shaded regions) result in failed slot allocation. The satellite sends a Variable Clear to Send (VCTS) message to confirm successful reservations. **Transmission Phase**: Nodes with confirmed reservations transmit their data in the allocated slots. Successfully transmitted data packets (from Nodes 1 and 2) mark a successful communication round. **Failed Rounds**: In subsequent reservation phases, collisions from overlapping RTS messages prevent successful data transmissions, leading to retries in future rounds. (From [Ort+21])

the throughput of DQ and 19 times the node energy efficiency of RESS-IoT while maintaining high fairness and significantly reducing protocol overhead. This makes MSDQ a highly effective solution for large-scale DtS-IoT networks, addressing scalability and energy constraints while ensuring robust channel access.

DRAAF The paper in [Ger+24a], also led by Jason Gerard, introduces the DtS-IoT Resource Allocation Analysis Framework (DRAAF), addressing the problem of evaluating and optimizing MAC protocols for DtS-IoT networks. While not a custom protocol, the proposed framework incorporates orbital mechanics, realistic energy consumption models, and characteristics of the LoRa-based physical layer. It uses a comprehensive set of performance metrics, such as throughput, delay, fairness, energy efficiency, and overhead. DRAAF is applied to two reservation-based MAC protocols: Reserve-and-Send Sift-IoT (RESS-IoT) and Distributed Queuing (DQ). Figure D.10 results are obtained from DRAAF. Results reveal that DQ outperforms RESS-IoT in throughput and node energy efficiency due to its efficient contention resolution strategy. At the same time, RESS-IoT achieves lower delays and matches DQ in fairness and satellite energy efficiency. The study highlights that despite RESS-IoT's higher overhead, it maintains competitive throughput through its reverse-exponential reservation mechanism. DRAAF's modularity and realistic scenario modeling make it a powerful tool for benchmarking and improving resource allocation mechanisms in DtS-IoT networks, with publicly available implementations fostering further research and protocol innovation.

D.2.1.5 Management

This paper [Sin+24] addresses the challenge of dynamic scheduling in NTN integrating LEO satellite constellations with terrestrial networks to optimize resource utilization and service continuity. This work is the first led by Arnau Singla, from i2CAT, which I started to co-supervise with Joan Ruiz-De-Azua. The proposed method enhances a Constellation Management System (CMS) by incorporating closed-loop feedback mechanisms using real-time telemetry data, such as satellite power and memory levels, to reschedule tasks dynamically. The CMS integrates scheduling and network management functionalities by generating adaptive task schedules for satellite operations, considering real-time feedback and failed task reallocation. The methodology involves a modular CMS architecture capable of handling telemetry updates, scenario changes, and reconfigurations, validated through rescheduling tests under simulated conditions. The results demonstrate a 10% improvement in throughput using the closed-loop approach compared to an open-loop baseline, highlighting the system's ability to adapt to fluctuations in telemetry and task failures rapidly. This research advances dynamic scheduling methodologies, enhancing the robustness and adaptability

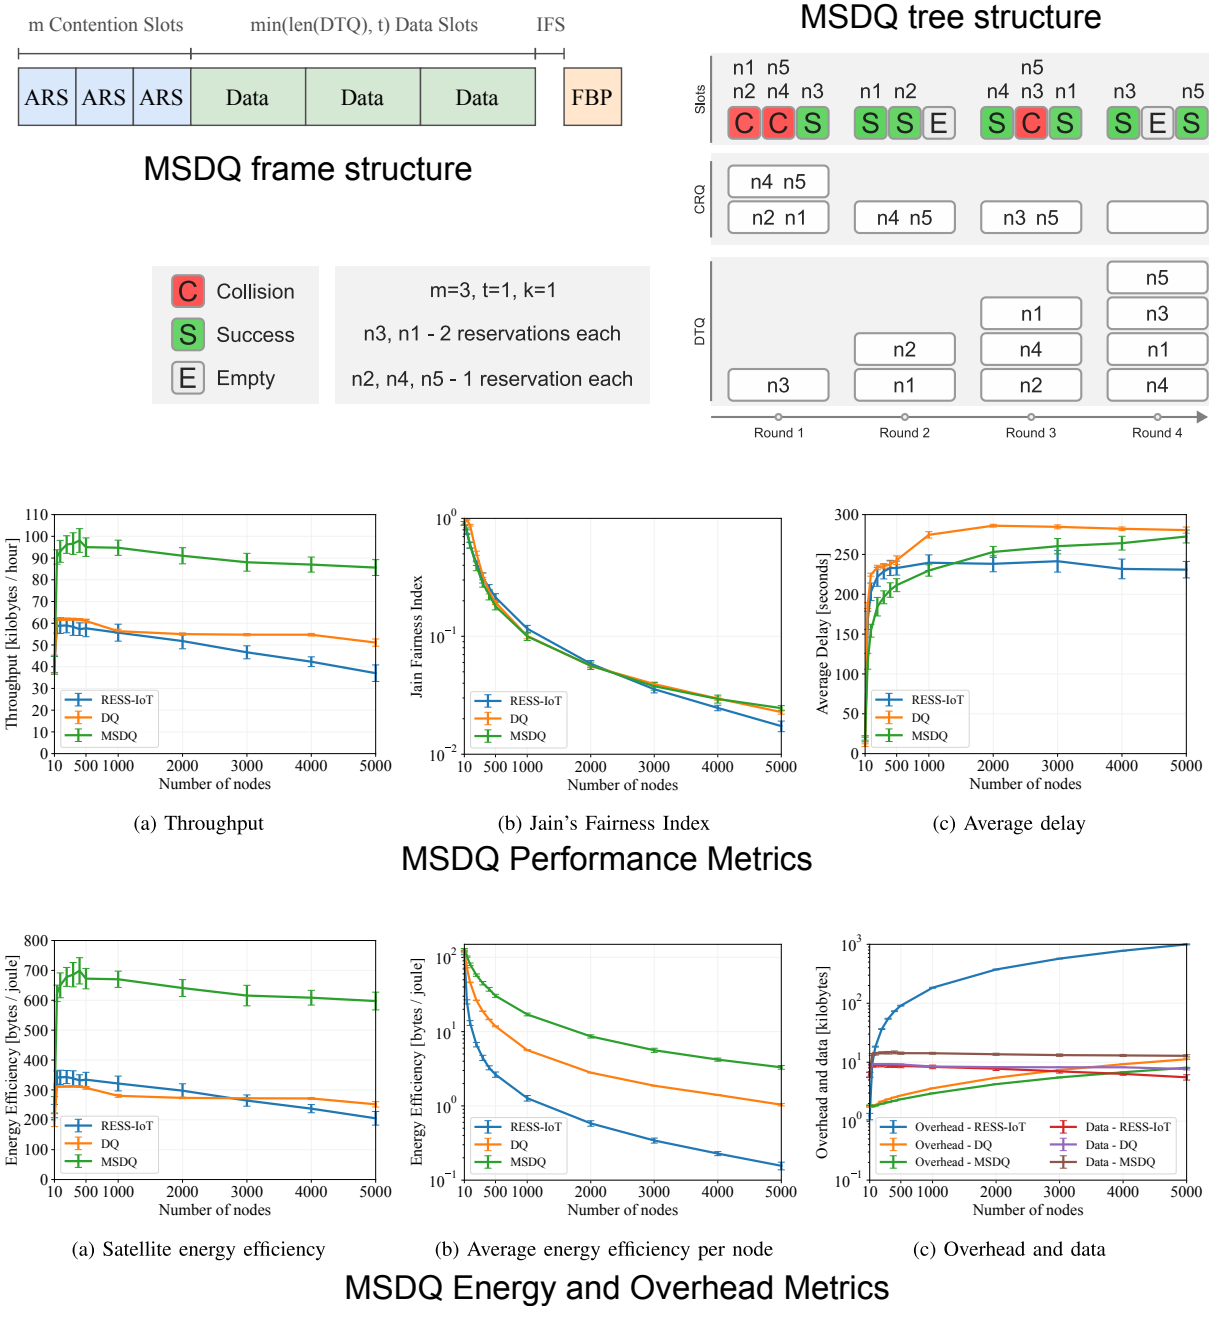

Figure D.10: Massively Scalable Distributed Queuing (MSDQ) protocol overview. **Top Left Panel**: The MSDQ frame structure, consisting of contention slots (ARS), data slots, and a final beacon packet (FBP), is designed to allocate resources efficiently. Contention results (Collision, Success, Empty) are managed within a dynamic tree structure to optimize reservations. **Top Right Panel**: MSDQ tree structure for collision resolution. Nodes progress through multiple contention rounds, dynamically branching to resolve collisions (C), allocate successful slots (S), and mark empty slots (E). The example shows how node reservations are distributed over four rounds. **Bottom Panels**: *Throughput*: MSDQ achieves significantly higher throughput. *Jain's Fairness Index*: MSDQ maintains fairness across nodes. *Average Delay*: MSDQ minimizes delay . *Satellite Energy Efficiency*: MSDQ maximizes energy efficiency at the satellite level. *Node Energy Efficiency*: MSDQ achieves superior average energy efficiency per node. *Overhead and Data Metrics*: MSDQ balances protocol overhead and transmits data effectively. (From [GFC24])

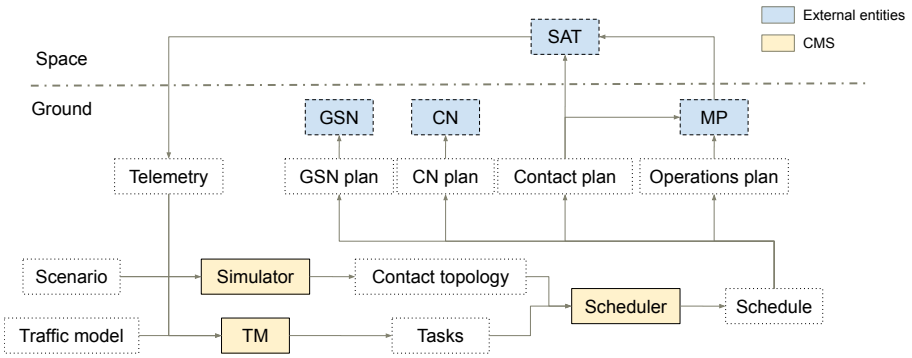

Figure D.11: Constellation Management System (CMS) for NTNs. The figure presents the CMS. It integrates external entities such as satellites (SAT), ground station networks (GSN), core networks (CN), and mission planners (MP). Satellites provide telemetry data, execute operations plans, and establish contact plans, while GSN and CN manage ground-to-satellite and core network connectivity through their plans. A simulator processes scenarios and traffic models, generating contact topology and network conditions to guide the scheduler. The scheduler integrates telemetry, topology, and task inputs to produce optimized schedules, while the traffic model defines task requirements and traffic patterns for accurate simulation. The CMS outputs detailed GSN, CN, contact, and operations plans external entities use to ensure efficient resource allocation. (From [Sin+24])

of NTNs for evolving 6G applications.

At the time of writing, several papers by Arnau Singla at i2CAT are under revision across various venues. His exceptional work on advancing the CMS and its integration with NB-IoT networks provides a unique opportunity to complement my LPWAN expertise, currently centered on LoRa/LoRaWAN technologies. Most relevantly, the CMS is scheduled to be flight validated in the recently deployed Sateliot satellites with NB-IoT payloads.

D.2.1.6 DTN Integration

The DtS-IoT and DTN domains are incredibly close, yet no integration has been materialized. This is a personal goal I intend to achieve soon. I'm likely in a unique position to do so.

DTN of Things My idea of integrating DTN and IoT is not new. Published as part of the book Delay and Disruption Tolerant Networks: Interplanetary and Earth-Bound — Architecture, Protocols, and Applications (CRC Press, 2018), I led a book chapter in [FF18] that introduces the concept of the *Delay-Tolerant Internet of Things*. This novel networking paradigm expands the IoT's capabilities to address intermittent connectivity and high-latency environments. Building upon the broader framework of DTN, I explored

the unique challenges and opportunities presented by integrating IoT into delay- and disruption-prone scenarios. The chapter provides a comprehensive discussion of the concept of a network of things, including its applications, architecture, and services, as well as an overview of existing IoT communication protocols and their limitations in such contexts. DTN of Things offers a model that decouples IoT connectivity from the constraints of continuous Internet-based communication. This enables scalable and resilient operations across diverse and resource-constrained environments.

DTN-based MAC Another exploration in the context of RF links was the consideration of DTN contact plans as a means of coordinating shared medium access across multiple satellites. The paper in [CFR18] addresses the integration challenges of TDMA-based Medium Access Control (MAC) with Scheduled DTN, focusing on the misalignment of communication times and the resulting inefficiencies in routing. The proposed method introduces three cross-layer approaches to incorporate TDMA slot information into CGR: (1) no knowledge of slotting, treating the channel as contention-prone and resulting in frequent rerouting; (2) partial knowledge, providing average contact data rates to improve volume estimates while maintaining low complexity; and (3) perfect knowledge, splitting contacts into detailed time slots for precise modeling, but at a higher computational cost. The approaches are tested on a simulated 16-node Walker Delta LEO constellation, evaluating performance under varying traffic loads. Key outcomes reveal that the partial knowledge approach achieves a balance of low computational complexity and high delivery rates, closely matching the performance of the perfect knowledge scheme at medium traffic levels. In contrast, while computationally expensive, the perfect knowledge scheme minimizes rerouting and delays. The results emphasize the importance of cross-layer information sharing in optimizing TDMA-based Scheduled DTNs for reliable and efficient data delivery.

CGR in DtS-IoT Within the STEREO project, and with Diego Maldonado 's participation, I supervised Sebastian Montoya during his master's internship to investigate the integration of CGR and DtS-IoT using the FLoRaSat simulation framework. The results were collected in the paper in [Mon+24]. The problem addressed involves optimizing data routing in DtS-IoT environments characterized by high latency, intermittent connectivity, and resource constraints. The study implements CGR algorithms within the FLoRaSat framework, integrating a contact plan generator and routing modules to simulate data flow through ISLs in Walker Delta and Star constellations. Key findings demonstrate that DTN's store-carry-and-forward approach, augmented by CGR, significantly reduces end-to-end delays. Key results and a depiction of the evaluated configurations are provided in Figure D.12. For instance, Star constellations equipped with ISLs attain delivery

delays as low as 10 minutes, while Delta constellations with eight orbital planes exhibit near-zero delays for fleets with over 24 satellites. The analysis also highlights the trade-offs in hop count and routing effort, showcasing the proposed integration's scalability and efficiency. These results underline the potential of DTN and CGR to revolutionize data delivery in DtS-IoT systems.

LoRa in DTN In the context of the MISSION project, Germán Scapin, a Master's student I co-supervised visited D3TN GmbH in Dresden, Germany, and developed an integration of LoRa with the μD3TN protocol stack. The outcome is summarized in the blog post titled Integrating LoRa with Delay-Tolerant Networking: A Practical Guide. It explores two methods to enhance IoT communications in environments with intermittent connectivity, such as rural areas and disaster recovery scenarios. The first method involves embedding μD3TN nodes between LoRa gateways and The Things Network (TTN), enabling a store-and-forward mechanism that ensures data persistence during connectivity interruptions. The second method utilizes DTN7's LoRa-specific CLA to encapsulate LoRa data into DTN bundles, facilitating reliable transmission and seamless interoperability between LoRa devices and μD3TN nodes. Both approaches demonstrated enhanced resilience, cross-protocol compatibility, and scalability, highlighting LoRa's potential role within modular DTN-based architectures and paving the way for future network expansion and adaptability.

D.2.2 Direct-to-Cell Access

The concept of Direct-to-Cell (D2C) refers to the direct connection between a standard handheld mobile device—beyond IoT-focused applications—and a LEO satellite. In the emerging community, the terms Direct-to-Device (D2D) and Direct-to-Smartphone (D2S) are often used interchangeably. Leveraging my expertise in DtS-IoT and my broader satellite communication background, I recently ventured into this area by hiring Hannah Pasandi as part of the STEREO project to advance research in this promising domain.

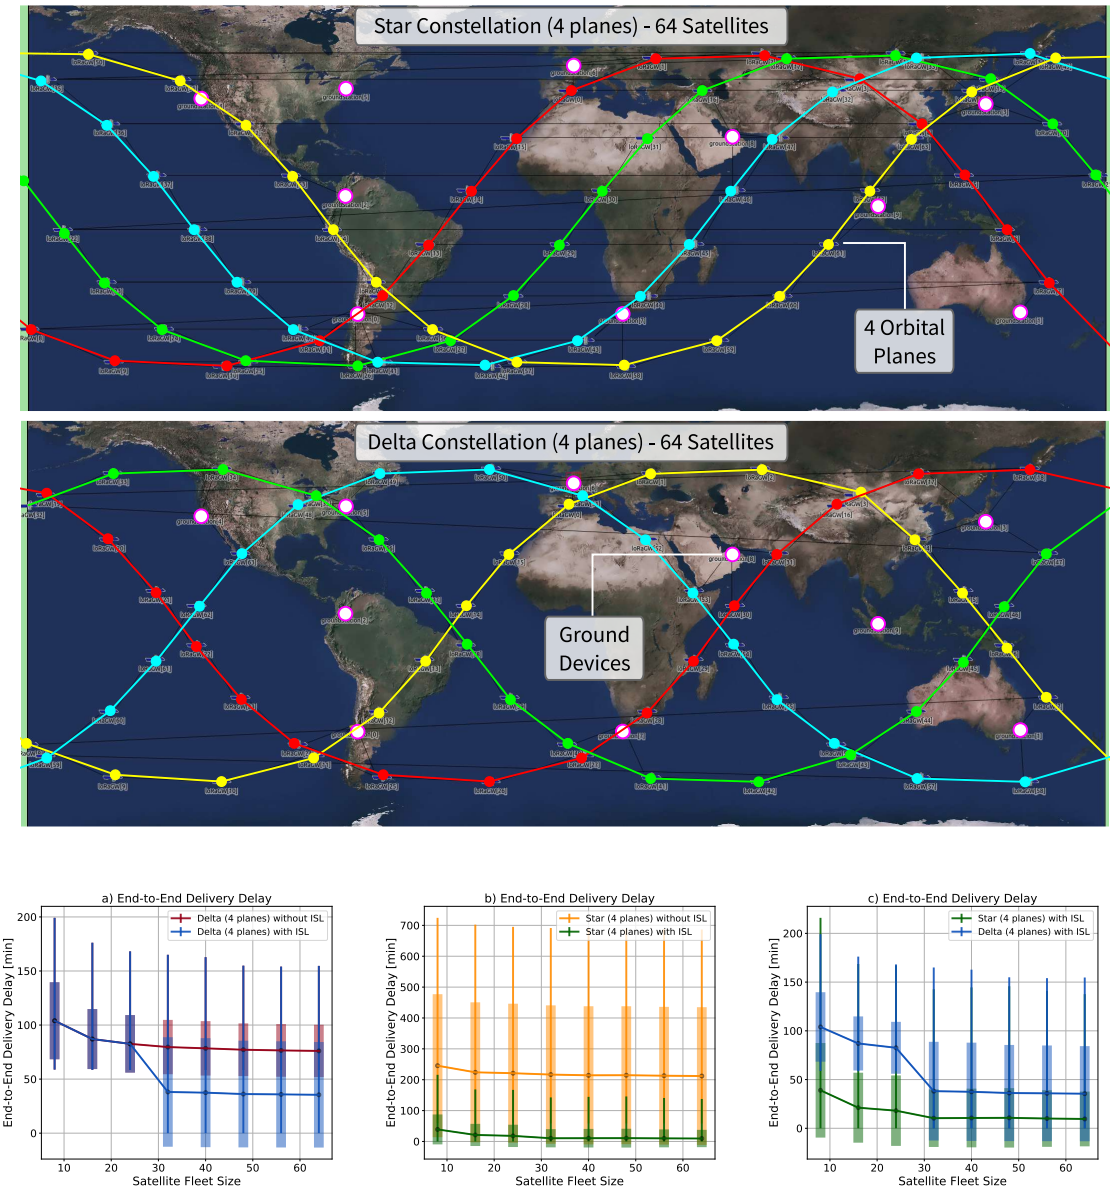

Figure D.12: Integration of CGR and DtS-IoT. **Top Panel**: Screenshots from the FLoRaSat simulator illustrating the orbital configurations of two satellite constellations with 64 satellites each. **Bottom Panel**: End-to-end delivery delay metrics for varying satellite fleet sizes in both constellations. Subfigure (a) demonstrates significant delay reductions when ISLs are utilized. Subfigure (b) reveals lower delays overall compared to the Delta constellation. Subfigure (c) highlights the Delta constellation's near-zero delays for larger fleets with ISLs, while the Star constellation maintains stable performance for all fleet sizes. Error bars indicate variability across different traffic patterns. (From [Mon+24])

The **Direct-to-Cell (D2C)** domain is poised to become a significant research niche, driven by several converging factors and with a distinct focus. Successfully scaling D2C connections would be a global game changer, enabling seamless service for all devices worldwide, potentially under a single operator. This transformative potential has spurred intense competition, with players like Starlink from SpaceX and others leading the race. However, the field is predominantly industry-driven, with limited academic research addressing its challenges. While academic efforts may trail behind industry advancements, I see immense potential in this domain as a fertile area for future research and innovation.

Our first paper in [Pas+24] comprehensively explores D2C satellite communication systems by addressing their technological, regulatory, and operational challenges. It identifies key issues, such as the need for robust protocols to handle the frequent handovers, high Doppler shifts, and variable link conditions inherent to LEO satellite constellations. To address these, the survey categorizes advancements in advanced beamforming techniques, software-defined payloads, and dynamic spectrum-sharing frameworks, highlighting their role in enhancing signal quality, link reliability, and spectral efficiency. The paper also underscores critical gaps in scalability for IoT-focused applications, energy-efficient protocol design, and spectrum coexistence strategies between terrestrial and satellite systems. Key outcomes include a detailed taxonomy of D2C systems, insights into current industry deployments, and proposed innovations such as adaptive bit-rate protocols and AI-driven spectrum management to address scalability and efficiency. This work positions D2C communication as a transformative technology for global connectivity, emphasizing its potential to bridge coverage gaps in remote and underserved regions.

The future of D2C depends on overcoming key challenges and leveraging emerging opportunities not only in technology, as discussed above, but also in business models and standardization. On the business front, scalable adoption requires cost-effective models leveraging synergies between Mobile Network Operators (MNOs) and Satellite Network Operators (SNOs), dynamic spectrum sharing, and tiered pricing structures, particularly in underserved regions. Standardization efforts, such as 3GPP radio interface harmonization and global spectrum management, are essential for seamless cross-network compatibility and stakeholder collaboration. Future directions should focus on interoperability frameworks and multi-tenant satellite access protocols to drive innovation and scalability in the D2C ecosystem. I consider this a significant challenge in the context of D2C.

Bibliography

[Alf+14] Marcos Alfonzo, Juan A Fraire, Esteban Kocian, and Nicolás Alvarez. “Development of a DTN bundle protocol convergence layer for SpaceWire”. In: *2014 IEEE Biennial Congress of Argentina (ARGENCON)*. IEEE. 2014, pp. 770–775. DOI: 10.1109/ARGENCON.2014.6868586. URL: https://doi.org/10.1109/ARGENCON.2014.6868586.

[Alv+21] Guido Alvarez, Mayco Dagatti, Matías Olmedo, Neyra Poveda, and Juan A Fraire. “Satellite-Based AIS Trade-off Analysis in the Context of the PANSAT CubeSat Mission”. In: *2021 IEEE 8th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2021, pp. 77–84. DOI: 10.1109/SMC-IT51442.2021.00017. URL: https://doi.org/10.1109/SMC-IT51442.2021.00017.

[Álv+22] Guido Álvarez, Juan A. Fraire, Khaled Abdelfadeel Hassan, Sandra Céspedes, and Dirk Pesch. “Uplink Transmission Policies for LoRa-Based Direct-to-Satellite IoT”. In: *IEEE Access* 10 (2022), pp. 72687–72701. DOI: 10.1109/ACCESS.2022.3189647. URL: https://doi.org/10.1109/ACCESS.2022.3189647.

[Ber+24] William Bernardoni, Robert Kassouf-Short, Robert Cardona, Brian Heller, Justin Curry, David Spivak, and Juan A Fraire. “Network Storage Analysis via Semiring Geometry”. In: *2024 IEEE Aerospace Conference*. IEEE. 2024, pp. 1–19. DOI: 10.1109/AERO58975.2024.10521207. URL: 10.1109/AERO58975.2024.10521207.

[Bul+24] Ronny L. Bull, Rachel M. Dudukovich, Juan A. Fraire, Nadia Kortas, Robert Kassouf-Short, Aaron Smith, and Ethan Schweinsberg. “Network Emulation Testbed Capabilities for Prototyping Space DTN Software and Protocols”. In: *IEEE INFOCOM 2024 - IEEE Conference on Computer Communications Workshops, Vancouver, BC, Canada, May 20, 2024*. IEEE, 2024, pp. 1–8. DOI: 10.1109/INFOCOMWKSHPS61880.2024.10620672. URL: https://doi.org/10.1109/INFOCOMWKSHPS61880.2024.10620672.

[Cao+18] Bin Cao, Ruhai Wang, Alaa Sabbagh, Siwei Peng, Kanglian Zhao, Juan A. Fraire, Guannan Yang, and Yue Wang. “Expected File-Delivery Time of DTN Protocol over Asymmetric Space Internetwork Channels”. In: *6th IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2018, Huntsville, AL, USA, December 11-13, 2018*. IEEE, 2018, pp. 147–151. DOI: 10.1109/WISEE.2018.8637352. URL: https://doi.org/10.1109/WiSEE.2018.8637352.

[Cap+22] Gabriel Maiolini Capez, Santiago Henn, Juan A. Fraire, and Roberto Garello. “Sparse Satellite Constellation Design for Global and Regional Direct-to-Satellite IoT Services”. In: *IEEE Trans. Aerosp. Electron. Syst.* 58.5 (2022), pp. 3786–3801. DOI: 10.1109/TAES.2022.3185970. URL: https://doi.org/10.1109/TAES.2022.3185970.

[Cap+24] Gabriel Maiolini Capez, Mauricio A Cáceres, Roberto Armellin, Chris P Bridges, Juan A Fraire, Stefan Frey, and Roberto Garello. “On the Use of Mega Constellation Services in Space: Integrating LEO Platforms into 6G Non-Terrestrial Networks”. In: *IEEE Journal on Selected Areas in Communications* (2024). DOI: 10.1109/JSAC.2024.3459078. URL: https://doi.org/10.1109/JSAC.2024.3459078.

[CFR18] Michael Carosino, Juan A. Fraire, and James A. Ritcey. “Integrating Scheduled DTNs and TDMA-based MAC Sublayers: Preliminary Results”. In: *6th IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2018, Huntsville, AL, USA, December 11-13, 2018*. IEEE, 2018, pp. 141–146. DOI: 10.1109/WISEE.2018.8637344. URL: https://doi.org/10.1109/WiSEE.2018.8637344.

[Che+23] Renato Cherini, Ramiro Detke, Juan A. Fraire, Pablo G. Madoery, and Jorge M. Finochietto. “Toward Deep Digital Contact Tracing: Opportunities and Challenges”. In: *IEEE Pervasive Comput.* 22.4 (2023), pp. 15–25. DOI: 10.1109/MPRV.2023.3320987. URL: https://doi.org/10.1109/MPRV.2023.3320987.

[Cho+24] Alexander Choquenaira-Florez, Juan A Fraire, Hannah B. Pasandi, and Hervé Rivano. “On the Role of Machine Learning in Satellite Internet of Things: A Survey of Techniques, Challenges, and Future Directions”. In: *Computer Networks* (2024), p. 111018. DOI: 10.1016/j.comnet.2024.111018. URL: https://doi.org/10.1016/j.comnet.2024.111018.

[Cho+25] Alexander Y. Choquenaira-Florez, Robin Ohs, Juan A. Fraire, and Hervé Rivano. “FLoRaSat 2: Simulating Cross-Linked Direct-to-Satellite IoT LEO Constellations”. In: *14th Advanced Satellite Multimedia Systems Conference*

and the 20th Signal Processing for Space Communications Workshop, ASM-S/SPSC 2025. In Press. IEEE, 2025.

[Cle+17] JA Clemente, JA Fraire, M Solinas, FJ Franco, Francesca Villa, Solenne Rey, Maud Baylac, Helmut Puchner, H Mecha, and Raoul Velazco. "SEU sensitivity trends of three successive generations of COTS SRAMs at ultra low bias voltage". In: *Conference on Radiation Effects on Components and Systems (RADECS 2017)*. 2017. URL: https://in2p3.hal.science/in2p3-01735635v1.

[Cle+18] Juan Antonio Clemente, Guillaume Hubert, Juan Fraire, Francisco J Franco, Francesca Villa, Solenne Rey, Maud Baylac, Helmut Puchner, Hortensia Mecha, and Raoul Velazco. "SEU characterization of three successive generations of COTS SRAMs at ultralow bias voltage to 14.2-MeV neutrons". In: *IEEE Transactions on Nuclear Science* 65.8 (2018), pp. 1858–1865. DOI: 10.1109/TNS.2018.2800905. URL: https://doi.org/10.1109/TNS.2018.2800905.

[Coe+16] A Coelho, R Laurent, E Mazer, Nacer-Eddine Zergainoh, J Fraire, M Solinas, S Karaoui, and Raoul Velazco. "Evidences of Bayesian Machines Robustness Against SEUs and SETs". In: *Radiation and its Effects on Components and Systems (RADECS'16)*. 2016. DOI: 10.1109/RADECS.2016.8093156. URL: https://doi.org/10.1109/RADECS.2016.8093156.

[Coe+17] Alexandre Coelho, Raphaël Laurent, Miguel Solinas, Juan Fraire, Emmanuel Mazer, Nacer-Eddine Zergainoh, Said Karaoui, and Raoul Velazco. "On the robustness of stochastic bayesian machines". In: *IEEE Transactions on Nuclear Science* 64.8 (2017), pp. 2276–2283. DOI: 10.1109/TNS.2017.2678204. URL: https://doi.org/10.1109/TNS.2017.2678204.

[Coe+18] Alexandre Coelho, Amir Charif, Nacer-Eddine Zergainoh, Juan A. Fraire, and Raoul Velazco. "A soft-error resilient route computation unit for 3D Networks-on-Chips". In: *2018 Design, Automation & Test in Europe Conference & Exhibition, DATE 2018, Dresden, Germany, March 19-23, 2018*. Ed. by Jan Madsen and Ayse K. Coskun. IEEE, 2018, pp. 1357–1362. DOI: 10.23919/DATE.2018.8342224. URL: https://doi.org/10.23919/DATE.2018.8342224.

[DAr+22] Pedro R. D'Argenio, Juan A. Fraire, Arnd Hartmanns, and Fernando D. Raverta. "Comparing Statistical and Analytical Routing Approaches for Delay-Tolerant Networks". In: *Quantitative Evaluation of Systems - 19th International Conference, QEST 2022, Warsaw, Poland, September 12-16, 2022, Proceedings*. Ed. by Erika Ábrahám and Marco Paolieri. Vol. 13479. Lecture Notes in

Computer Science. Springer, 2022, pp. 337–355. DOI: 10.1007/978-3-031-16336-4_17. URL: https://doi.org/10.1007/978-3-031-16336-4%5C_17.

[DAr+24] Pedro R D'Argenio, Juan Fraire, Arnd Hartmanns, and Fernando Raverta. "Comparing Statistical, Analytical, and Learning-Based Routing Approaches for Delay-Tolerant Networks". In: *ACM Transactions on Modeling and Computer Simulation* (2024). DOI: 10.1145/3665927. URL: https://doi.org/10.1145/3665927.

[DF20] Olivier De Jonckère and Juan A Fraire. "A shortest-path tree approach for routing in space networks". In: *China Communications* 17.7 (2020), pp. 52–66. DOI: 10.23919/J.CC.2020.07.005. URL: https://doi.org/10.23919/J.CC.2020.07.005.

[DFB24] Olivier De Jonckere, Juan A Fraire, and Scott Burleigh. "Distributed Volume Management in Space DTNs: Scoping Schedule-Aware Bundle Routing". In: *2024 IEEE 10th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2024, pp. 125–134. DOI: 10.1109/SMC-IT61443.2024.00021. URL: https://doi.org/10.1109/SMC-IT61443.2024.00021.

[DFH20] Pedro R. D'Argenio, Juan A. Fraire, and Arnd Hartmanns. "Sampling Distributed Schedulers for Resilient Space Communication". In: *NASA Formal Methods - 12th International Symposium, NFM 2020, Moffett Field, CA, USA, May 11-15, 2020, Proceedings*. Ed. by Ritchie Lee, Susmit Jha, and Anastasia Mavridou. Vol. 12229. Lecture Notes in Computer Science. Springer, 2020, pp. 291–310. DOI: 10.1007/978-3-030-55754-6_17. URL: https://doi.org/10.1007/978-3-030-55754-6_17.

[Dow+20] Alexis Alfredo Dowhuszko, Juan A. Fraire, Musbah Shaat, and Ana I. Pérez-Neira. "LEO satellite constellations to offload optical terrestrial networks in placement of popular content in 5G edge nodes". In: *22nd International Conference on Transparent Optical Networks, ICTON 2020, Bari, Italy, July 19-23, 2020*. IEEE, 2020, pp. 1–6. DOI: 10.1109/ICTON51198.2020.9203447. URL: https://doi.org/10.1109/ICTON51198.2020.9203447.

[FCA19] Juan A. Fraire, Sandra Céspedes, and Nicola Accettura. "Direct-To-Satellite IoT - A Survey of the State of the Art and Future Research Perspectives - Backhauling the IoT Through LEO Satellites". In: *Ad-Hoc, Mobile, and Wireless Networks - 18th International Conference on Ad-Hoc Networks and Wireless, ADHOC-NOW 2019, Luxembourg, Luxembourg, October 1-3, 2019, Proceedings*. Ed. by Maria Rita Palattella, Stefano Scanzio, and Sinem Co-

leri Ergen. Vol. 11803. Lecture Notes in Computer Science. Springer, 2019, pp. 241–258. DOI: 10.1007/978-3-030-31831-4_17. URL: https://doi.org/10.1007/978-3-030-31831-4%5C_17.

[FD21] Juan A. Fraire and Juan Eduardo Durán. "Revising Computer Science Networking Hands-On Courses in the Context of the Future Internet". In: *IEEE Trans. Educ.* 64.2 (2021), pp. 133–138. DOI: 10.1109/TE.2020.3015673. URL: https://doi.org/10.1109/TE.2020.3015673.

[Fel+22] Marius Feldmann, Juan A. Fraire, Felix Walter, and Scott C. Burleigh. "Ring Road Networks: Access for Anyone". In: *IEEE Commun. Mag.* 60.4 (2022), pp. 38–44. DOI: 10.1109/MCOM.001.2100835. URL: https://doi.org/10.1109/MCOM.001.2100835.

[Fer+19] Pablo A. Ferreyra, Juan A. Fraire, Fabian Gomez, Raoul Velazco, Daniel Sánchez, and Dardo Vinas Viscardi. "Delay-Tolerant Wireless Networks on Chip: Preliminary Analysis and Results". In: *IEEE Latin American Test Symposium, LATS 2019, Santiago, Chile, March 11-13, 2019*. IEEE, 2019, pp. 1–6. DOI: 10.1109/LATW.2019.8704623. URL: https://doi.org/10.1109/LATW.2019.8704623.

[Fer+21] Pablo A. Ferreyra, Rubén Danilo Capkob, Alberto Fabián Gómez, Juan Andres Fraire, and Carlos José Barrientos. "Embedded wireless delay tolerant networks on chips for segmented architectures". In: *Int. J. Embed. Syst.* 14.6 (2021), pp. 578–591. DOI: 10.1504/IJES.2021.121087. URL: https://doi.org/10.1504/IJES.2021.121087.

[Fer+24] Elena Fernández-Niño, Juan A. Fraire, Sergi Figuerola, Adriano Camps, and Joan Adrià Ruiz De Azua. "A Case Study of an Hybrid RF and Optical Inter-Satellite Link Terminal to Enhance Optical Pointing". In: *Joint European Conference on Networks and Communications & 6G Summit, EuCNC/6G Summit 2024, Antwerp, Belgium, June 3-6, 2024*. IEEE, 2024, pp. 640–645. DOI: 10.1109/EUCNC/6GSUMMIT60053.2024.10597073. URL: https://doi.org/10.1109/EuCNC/6GSummit60053.2024.10597073.

[FF15a] Juan A. Fraire and Jorge M. Finochietto. "Design challenges in contact plans for disruption-tolerant satellite networks". In: *IEEE Commun. Mag.* 53.5 (2015), pp. 163–169. DOI: 10.1109/MCOM.2015.7105656. URL: https://doi.org/10.1109/MCOM.2015.7105656.

[FF15b] Juan A. Fraire and Jorge M. Finochietto. "Routing-aware fair contact plan design for predictable delay tolerant networks". In: *Ad Hoc Networks* 25 (2015), pp. 303–313. DOI: 10.1016/J.ADHOC.2014.07.006. URL: https://doi.org/10.1016/j.adhoc.2014.07.006.

[FF18] Juan A Fraire and Jorge M Finochietto. “DTN of Things”. In: *Delay and Disruption Tolerant Networks*. CRC Press, 2018, pp. 275–308.

[FFB17a] Juan A Fraire, Jorge M Finochietto, and Scott C Burleigh. *Delay Tolerant Satellite Networks*. Artech House, 2017. ISBN: 9781630815172.

[FFB17b] Juan A. Fraire, Marius Feldmann, and Scott C. Burleigh. “Benefits and challenges of cross-linked ring road satellite networks: A case study”. In: *IEEE International Conference on Communications, ICC 2017, Paris, France, May 21-25, 2017*. IEEE, 2017, pp. 1–7. DOI: 10.1109/ICC.2017.7996778. URL: https://doi.org/10.1109/ICC.2017.7996778.

[FFM13a] Juan A Fraire, Alejandro Ferreyra, and Carlos Marques. “OpenCL overview, implementation, and performance comparison”. In: *IEEE Latin America Transactions* 11.1 (2013), pp. 274–280. DOI: 10.1109/TLA.2013.6502816. URL: https://doi.org/10.1109/TLA.2013.6502816.

[FFM13b] Juan Andrés Fraire, Pablo Ferreyra, and Carlos Marques. “Opencl-accelerated simplified general perturbations 4 algorithm”. In: *XIV Argentine Symposium on Technology (AST)-JAIIO 42 (2013)*. 2013. URL: http://sedici.unlp.edu.ar/handle/10915/76850.

[FFW18] Marius Feldmann, Juan A. Fraire, and Felix Walter. “Tracking Lunar Ring Road Communication”. In: *2018 IEEE International Conference on Communications, ICC 2018, Kansas City, MO, USA, May 20-24, 2018*. IEEE, 2018, pp. 1–7. DOI: 10.1109/ICC.2018.8423031. URL: https://doi.org/10.1109/ICC.2018.8423031.

[FG21] Juan A. Fraire and Elías L. Gasparini. “Centralized and Decentralized Routing Solutions for Present and Future Space Information Networks”. In: *IEEE Netw.* 35.4 (2021), pp. 110–117. DOI: 10.1109/MNET.011.2100102. URL: https://doi.org/10.1109/MNET.011.2100102.

[FGI23a] Juan Andres Fraire, Alexandre Guitton, and Oana Iova. “Recovering Headerless Frames in LR-FHSS”. In: *Proceedings of the 2023 International Conference on embedded Wireless Systems and Networks, EWSN 2023, Rende, Italy, September 25-27, 2023*. Ed. by Giancarlo Fortino, Valeria Loscrı`, Fabrizio Granelli, Tarek F. Abdelzaher, Antonella Guzzo, and Claudio Savaglio. ACM, 2023, pp. 170–180. DOI: 10.5555/3639940.3639963. URL: https://dl.acm.org/doi/10.5555/3639940.3639963.

[FGI23b] Juan Andrés A Fraire, Alexandre Guitton, and Oana Iova. “LR-FHSS: j’ai perdu l’entête”. In: *CoRes 2023-8èmes Rencontres Francophones sur la Conception de protocoles, l’évaluation de performances et l’expérimentation de*

Réseaux de communication. 2023. URL: https://hal.science/hal-04087040v1.

[FIV22] Juan A. Fraire, Oana Iova, and Fabrice Valois. "Space-Terrestrial Integrated Internet of Things: Challenges and Opportunities". In: *IEEE Commun. Mag.* 60.12 (2022), pp. 64–70. DOI: 10.1109/MCOM.008.2200215. URL: https://doi.org/10.1109/MCOM.008.2200215.

[FJB21] Juan A. Fraire, Olivier De Jonckère, and Scott C. Burleigh. "Routing in the Space Internet: A contact graph routing tutorial". In: *J. Netw. Comput. Appl.* 174 (2021), p. 102884. DOI: 10.1016/J.JNCA.2020.102884. URL: https://doi.org/10.1016/j.jnca.2020.102884.

[FK15] Juan A Fraire and Esteban G Kocian. "A Flexible and Autonomous Communication Subsystem for Constellations of Small-Satellites". In: *SSC15-P-9, 2015 AIAA/USU Conference on Small Satellites, Logan, Utah, USA* (2015). URL: https://digitalcommons.usu.edu/smallsat/2015/all2015/86/.

[FMF13] Juan A. Fraire, Pablo G. Madoery, and Jorge M. Finochietto. "On the design of fair contact plans in predictable Delay-Tolerant Networks". In: *IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2013, Baltimore, MD, USA, November 7-9, 2013*. IEEE, 2013, pp. 1–7. DOI: 10.1109/WISEE.2013.6737548. URL: https://doi.org/10.1109/WiSEE.2013.6737548.

[FMF14a] Juan A. Fraire, Pablo G Madoery, and Jorge M Finochietto. "On the design and analysis of fair contact plans in predictable delay-tolerant networks". In: *IEEE Sensors Journal* 14.11 (2014), pp. 3874–3882. DOI: 10.1109/JSEN.2014.2348917. URL: https://doi.org/10.1109/JSEN.2014.2348917.

[FMF14b] Juan A. Fraire, Pablo G. Madoery, and Jorge M. Finochietto. "Leveraging routing performance and congestion avoidance in predictable delay tolerant networks". In: *2014 IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2014, Noordwijk, Netherlands, October 30-31, 2014*. IEEE, 2014, pp. 1–7. DOI: 10.1109/WISEE.2014.6973079. URL: https://doi.org/10.1109/WiSEE.2014.6973079.

[FMF16] Juan A. Fraire, Pablo G. Madoery, and Jorge M. Finochietto. "Traffic-aware contact plan design for disruption-tolerant space sensor networks". In: *Ad Hoc Networks* 47 (2016), pp. 41–52. DOI: 10.1016/J.ADHOC.2016.04.007. URL: https://doi.org/10.1016/j.adhoc.2016.04.007.

[Fra+15a] Juan A. Fraire, Pablo G. Madoery, Jorge M. Finochietto, and Edward J. Birrane. “Congestion modeling and management techniques for predictable disruption tolerant networks”. In: *40th IEEE Conference on Local Computer Networks, LCN 2015, Clearwater Beach, FL, USA, October 26-29, 2015*. Ed. by Salil S. Kanhere, Jens Tölle, and Soumaya Cherkaoui. IEEE Computer Society, 2015, pp. 544–551. DOI: 10.1109/LCN.2015.7366369. URL: https://doi.org/10.1109/LCN.2015.7366369.

[Fra+15b] Juan A. Fraire, Pablo G. Madoery, Jorge M. Finochietto, and Guillermo Leguizamón. “Preliminary results of an evolutionary approach towards Contact Plan design for satellite DTNs”. In: *2015 IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2015, Orlando, FL, USA, December 14-16, 2015*. Ed. by Ron Brown and Ali Abedi. IEEE, 2015, pp. 1–7. DOI: 10.1109/WISEE.2015.7393097. URL: https://doi.org/10.1109/WiSEE.2015.7393097.

[Fra+16] Juan A. Fraire, Pablo G. Madoery, Jorge M. Finochietto, Pablo A. Ferreyra, and Raoul Velazco. “Internetworking approaches towards along-track segmented satellite architectures”. In: *2016 IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2016, Aachen, Germany, September 26-28, 2016*. IEEE, 2016, pp. 123–128. DOI: 10.1109/WISEE.2016.7877316. URL: https://doi.org/10.1109/WiSEE.2016.7877316.

[Fra+17a] Juan A Fraire, Pablo Madoery, Fernando Raverta, Jorge M Finochietto, and Raoul Velazco. “Dtnsim: Bridging the gap between simulation and implementation of space-terrestrial dtns”. In: *2017 6th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2017, pp. 120–123. DOI: 10.1109/SMC-IT.2017.27. URL: https://doi.org/10.1109/SMC-IT.2017.27.

[Fra+17b] Juan A. Fraire, Pablo G. Madoery, Scott C. Burleigh, Marius Feldmann, Jorge M. Finochietto, Amir Charif, Nacer-Eddine Zergainoh, and Raoul Velazco. “Assessing Contact Graph Routing Performance and Reliability in Distributed Satellite Constellations”. In: *J. Comput. Networks Commun.* 2017 (2017), 2830542:1–2830542:18. DOI: 10.1155/2017/2830542. URL: https://doi.org/10.1155/2017/2830542.

[Fra+17c] Juan A. Fraire, Pablo G. Madoery, Jorge M. Finochietto, and Guillermo Leguizamón. “An evolutionary approach towards contact plan design for disruption-tolerant satellite networks”. In: *Appl. Soft Comput.* 52 (2017), pp. 446–456. DOI: 10.1016/J.ASOC.2016.10.023. URL: https://doi.org/10.1016/j.asoc.2016.10.023.

[Fra+18a] Juan A. Fraire, Pablo G. Madoery, Amir Charif, and Jorge M. Finochietto. “On route table computation strategies in Delay-Tolerant Satellite Networks”. In: *Ad Hoc Networks* 80 (2018), pp. 31–40. DOI: 10.1016/J.ADHOC.2018.07.002. URL: https://doi.org/10.1016/j.adhoc.2018.07.002.

[Fra+18b] Juan A. Fraire, Gilles Nies, Holger Hermanns, Kristian Bay, and Morten Bisgaard. “Battery-Aware Contact Plan Design for LEO Satellite Constellations: The Ulloriaq Case Study”. In: *IEEE Global Communications Conference, GLOBECOM 2018, Abu Dhabi, United Arab Emirates, December 9-13, 2018*. IEEE, 2018, pp. 1–7. DOI: 10.1109/GLOCOM.2018.8647822. URL: https://doi.org/10.1109/GLOCOM.2018.8647822.

[Fra+19] Juan A. Fraire, Marius Feldmann, Felix Walter, Elena Fantino, and Scott C. Burleigh. “Networking in Interstellar Dimensions: Communicating With TRAPPIST-1”. In: *IEEE Trans. Aerosp. Electron. Syst.* 55.4 (2019), pp. 1656–1665. DOI: 10.1109/TAES.2018.2874149. URL: https://doi.org/10.1109/TAES.2018.2874149.

[Fra+20a] Juan A. Fraire, Santiago Henn, Fabio Dovis, Roberto Garello, and Giorgio Taricco. “Sparse Satellite Constellation Design for LoRa-based Direct-to-Satellite Internet of Things”. In: *IEEE Global Communications Conference, GLOBECOM 2020, Virtual Event, Taiwan, December 7-11, 2020*. IEEE, 2020, pp. 1–6. DOI: 10.1109/GLOBECOM42002.2020.9348042. URL: https://doi.org/10.1109/GLOBECOM42002.2020.9348042.

[Fra+20b] Juan A. Fraire, Gilles Nies, Carsten Gerstacker, Holger Hermanns, Kristian Bay, and Morten Bisgaard. “Battery-Aware Contact Plan Design for LEO Satellite Constellations: The Ulloriaq Case Study”. In: *IEEE Trans. Green Commun. Netw.* 4.1 (2020), pp. 236–245. DOI: 10.1109/TGCN.2019.2954166. URL: https://doi.org/10.1109/TGCN.2019.2954166.

[Fra+21] Juan A. Fraire, Carsten Gerstacker, Holger Hermanns, Gilles Nies, Morten Bisgaard, and Kristian Bay. “On the scalability of battery-aware contact plan design for LEO satellite constellations”. In: *Int. J. Satell. Commun. Netw.* 39.2 (2021), pp. 193–204. DOI: 10.1002/SAT.1374. URL: https://doi.org/10.1002/sat.1374.

[Fra+22a] Juan A. Fraire, Pablo G. Madoery, Mehdi Ait Mesbah, Oana Iova, and Fabrice Valois. “Simulating LoRa-Based Direct-to-Satellite IoT Networks with FLoRaSaT”. In: *23rd IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks, WoWMoM 2022, Belfast, United Kingdom, June 14-17, 2022*. IEEE, 2022, pp. 464–470. DOI: 10.1109/WOWMOM54355.2022.00072. URL: https://doi.org/10.1109/WoWMoM54355.2022.00072.

[Fra+22b] Juan Andrés A Fraire, Mehdi Mesbah, Pablo Madoery, Oana Iova, and Fabrice Valois. “FLoRaSat: Le simulateur pour la communication LoRa avec les satellites LEO”. In: *CORES 2022–7ème Rencontres Francophones sur la Conception de Protocoles, l’Évaluation de Performance et l’Expérimentation des Réseaux de Communication*. 2022, pp. 1–4. URL: https://hal.science/hal-03657890v1.

[Fra+24] Juan A. Fraire, Santiago Henn, Gregory Stock, Robin Ohs, Holger Hermanns, Felix Walter, Lynn Van Broock, Gabriel Ruffini, Federico Machado, Pablo Serratti, and Jose Relloso. “Quantitative analysis of segmented satellite network architectures: A maritime surveillance case study”. In: *Comput. Networks* 255 (2024), p. 110874. DOI: 10.1016/J.COMNET.2024.110874. URL: https://doi.org/10.1016/j.comnet.2024.110874.

[Fra17] Juan A Fraire. “Introducing contact plan designer: A planning tool for DTN-based space-terrestrial networks”. In: *2017 6th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2017, pp. 124–127. DOI: 10.1109/SMC-IT.2017.28. URL: https://doi.org/10.1109/SMC-IT.2017.28.

[FRF24] Juan A Fraire, Joan A Ruiz-de-Azua, and Elena Fernandez-Nino. “Towards Optimal Placement of Free-Space Optical Terminal on the Spacecraft Body”. In: *IEEE Transactions on Aerospace and Electronic Systems* (2024). DOI: 10.1109/TAES.2024.3435176. URL: https://doi.org/10.1109/TAES.2024.3435176.

[Ger+24a] Jason Gerard, Juan A Fraire, Raydel Ortigueira, and Sandra Céspedes. “DtS-IoT Resource Allocation Analysis Framework: Assessing DQ and RESS-IoT”. In: *2024 IEEE 10th World Forum on Internet of Things (WF-IoT)*. IEEE. 2024, pp. 672–677. DOI: 10.1109/WF-IoT62078.2024.10811213. URL: https://doi.org/10.1109/WF-IoT62078.2024.10811213.

[Ger+24b] Jason Gerard, Andre Ibrahim, Juan A. Fraire, and Sandra Céspedes. “Laser Link Scheduler: Autonomous Scheduling for Deep-Space Free-Space Optical Communications”. In: *IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. In Press. Daytona Beach, FL, USA: IEEE, Dec. 2024.

[GFC24] Jason Gerard, Juan A Fraire, and Sandra Céspedes. “Unlocking DtS-IoT Medium Access Through the Massively Scalable Distributed Queuing Protocol”. In: *2024 IEEE 10th World Forum on Internet of Things (WF-IoT)*. IEEE. 2024, pp. 840–845. DOI: 10.1109/WF-IoT62078.2024.10811352. URL: https://doi.org/10.1109/WF-IoT62078.2024.10811352.

[Gon+20] Shulei Gong, Hong Shen, Kanglian Zhao, Ruhai Wang, Xinggan Zhang, Tomaso de Cola, and Juan A. Fraire. “Network Availability Maximization for Free-Space Optical Satellite Communications”. In: *IEEE Wirel. Commun. Lett.* 9.3 (2020), pp. 411–415. DOI: 10.1109/LWC.2019.2958804. URL: https://doi.org/10.1109/LWC.2019.2958804.

[Hen+19] Santiago M Henn, Juan A Fraire, Pablo C Marino, Maximiliano Fischer, Andrea Bianco, and Holger Hermanns. “Towards the design of space missions with optimal accessibility via relay satellites”. In: *2019 IEEE International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2019, pp. 100–105. DOI: 10.1109/SMC-IT.2019.00018. URL: https://doi.org/10.1109/SMC-IT.2019.00018.

[HFH23] Santiago Henn, Juan A. Fraire, and Holger Hermanns. “Polygon-Based Algorithms for N-Satellite Constellations Coverage Computing”. In: *IEEE Trans. Aerosp. Electron. Syst.* 59.5 (2023), pp. 7166–7182. DOI: 10.1109/TAES.2023.3289479. URL: https://doi.org/10.1109/TAES.2023.3289479.

[Ila+25] Pablo Ilabaca, Diego Maldonado, Juan A. Fraire, Samuel Montejo-Sanchez, and Sandra Céspedes. “Constellation LoRa Optimistic Collision Information (CL-OCI): Scalable Network Size Estimation for DtS-IoT”. In: *Proceedings of the 2025 IEEE International Conference on Communications (ICC)*. In Press. Montreal, Canada: IEEE, June 2025, pp. 8–12.

[JF23] Olivier De Jonckère and Juan A. Fraire. “Inter-Regional Routing in Interplanetary Networks with Shortcuts and Contact Passageways”. In: *IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2023, Aveiro, Portugal, September 6-8, 2023*. IEEE, 2023, pp. 87–92. DOI: 10.1109/WISEE58383.2023.10289358. URL: https://doi.org/10.1109/WiSEE58383.2023.10289358.

[JFB23a] Olivier De Jonckère, Juan A. Fraire, and Scott Burleigh. “Enhanced Pathfinding and Scalability with Shortest-Path Tree Routing for Space Networks”. In: *IEEE International Conference on Communications, ICC 2023, Rome, Italy, May 28 - June 1, 2023*. IEEE, 2023, pp. 4082–4088. DOI: 10.1109/ICC45041.2023.10279514. URL: https://doi.org/10.1109/ICC45041.2023.10279514.

[JFB23b] Olivier De Jonckère, Juan A. Fraire, and Scott Burleigh. “On the Tractability of Yen’s Algorithm and Contact Graph Modeling in Contact Graph Routing”. In: *IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2023, Aveiro, Portugal, September 6-8, 2023*. IEEE, 2023,

pp. 80–86. DOI: 10.1109/WISEE58383.2023.10289594. URL: https://doi.org/10.1109/WiSEE58383.2023.10289594.

[Le +24] Alice Le Bihan, Juan A Fraire, Pierre Francois, and Felix Flentge. “IPN-V: The Interplanetary Network Visualizer”. In: *2024 IEEE 10th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2024, pp. 142–150. DOI: 10.1109/SMC-IT61443.2024.00023. URL: https://doi.org/10.1109/SMC-IT61443.2024.00023.

[Led+24] Oscar Ledesma, Paula Lamo, Juan A Fraire, María Ruiz, and Miguel A Sánchez. “Architectural framework and feasibility of internet of things-driven mars exploration via satellite constellations”. In: *Electronics* 13.7 (2024), p. 1289. DOI: 10.3390/electronics13071289. URL: https://doi.org/10.3390/electronics13071289.

[Liu+19] Lu Liu, Ruhai Wang, Yan Zhou, Wenfeng Li, Yi Jin, and Juan A Fraire. “Energy consumption of bundle protocol for reliable file delivery over deep-space channels”. In: *2019 IEEE International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2019, pp. 106–111. DOI: 10.1109/SMC-IT.2019.00019. URL: https://doi.org/10.1109/SMC-IT.2019.00019.

[LLF24] Oscar Ledesma, Paula Lamo, and Juan A Fraire. “Trends in LPWAN Technologies for LEO Satellite Constellations in the NewSpace Context”. In: *Electronics* 13.3 (2024), p. 579. DOI: 10.3390/electronics13030579. URL: https://doi.org/10.3390/electronics13030579.

[Mad+18a] Pablo G. Madoery, Juan A. Fraire, Fernando D. Raverta, and Jorge M. Finochietto. “On the Trade-off Analysis of Centralized and Distributed Routing in Delay Tolerant Satellite Networks”. In: *Proceedings of the IV School on Systems and Networks, SSN 2018, Valdivia, Chile, October 29-31, 2018*. Ed. by Sandra Céspedes and Javier Bustos-Jiménez. Vol. 2178. CEUR Workshop Proceedings. CEUR-WS.org, 2018, pp. 89–93. URL: https://ceur-ws.org/Vol-2178/SSN2018%5C_paper%5C_32.pdf.

[Mad+18b] Pablo G. Madoery, Juan A. Fraire, Fernando D. Raverta, Jorge M. Finochietto, and Scott C. Burleigh. “Managing Routing Scalability in Space DTNs”. In: *6th IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2018, Huntsville, AL, USA, December 11-13, 2018*. IEEE, 2018, pp. 177–182. DOI: 10.1109/WISEE.2018.8637324. URL: https://doi.org/10.1109/WiSEE.2018.8637324.

[Mad+18c] Pablo G. Madoery, Fernando D. Raverta, Juan A. Fraire, and Jorge M. Finochietto. “Routing in Space Delay Tolerant Networks under Uncertain Contact Plans”. In: *2018 IEEE International Conference on Communications, ICC 2018, Kansas City, MO, USA, May 20-24, 2018*. IEEE, 2018, pp. 1–6. DOI: 10.1109/ICC.2018.8422917. URL: https://doi.org/10.1109/ICC.2018.8422917.

[Mad+21] Pablo G. Madoery, Ramiro Detke, Lucas Blanco, Sandro Comerci, Juan A. Fraire, Aldana M. González-Montoro, Juan Carlos Bellassai, Grisel Maribel Britos, Silvia María Ojeda, and Jorge M. Finochietto. “Feature selection for proximity estimation in COVID-19 contact tracing apps based on Bluetooth Low Energy (BLE)”. In: *Pervasive Mob. Comput.* 77 (2021), p. 101474. DOI: 10.1016/J.PMCJ.2021.101474. URL: https://doi.org/10.1016/j.pmcj.2021.101474.

[Mad+24] Pablo G. Madoery, Juan A. Fraire, Jorge M. Finochietto, Halim Yanikomeroglu, and Gunes Karabulut Kurt. “A Novel Non- Terrestrial Networks Architecture: All Optical LEO Constellations with High-Altitude Ground Stations”. In: *IEEE International Conference on Communications Workshops, ICC 2024 Workshops, Denver, CO, USA, June 9-13, 2024*. IEEE, 2024, pp. 1115–1120. DOI: 10.1109/ICCWORKSHOPS59551.2024.10615868. URL: https://doi.org/10.1109/ICCWorkshops59551.2024.10615868.

[Mal+23] Diego Maldonado, Juan A Fraire, Pablo Ilabaca, Hervé Rivano, and Sandra Céspedes. “Network Size Estimation for LoRa-Based Direct-to-Satellite IoT”. In: *2023 IEEE Cognitive Communications for Aerospace Applications Workshop (CCAAW)*. IEEE. 2023, pp. 1–6. DOI: 10.1109/CCAAW57883.2023.10219363. URL: https://doi.org/10.1109/CCAAW57883.2023.10219363.

[Mal+24a] Diego Maldonado, Leonardo S Cardoso, Juan A Fraire, Alexandre Guitton, Oana Iova, Megumi Kaneko, and Hervé Rivano. “Enhanced LR-FHSS receiver for headerless frame recovery in space-terrestrial integrated IoT networks”. In: *Computer Networks* (2024), p. 111018. DOI: 10.1016/j.comnet.2024.111018. URL: https://doi.org/10.1016/j.comnet.2024.111018.

[Mal+24b] Diego Maldonado, Megumi Kaneko, Juan A. Fraire, Alexandre Guitton, Oana Iova, and Hervé Rivano. “Enhancing LR-FHSS Scalability Through Advanced Sequence Design and Demodulator Allocation”. In: *CoRR* abs/2407.03490 (2024). Under Evaluation for IEEE TGCN. DOI: 10.48550/ARXIV.2407.03490. arXiv: 2407.03490. URL: https://doi.org/10.48550/arXiv.2407.03490.

[MFF16] Pablo G Madoery, Juan A Fraire, and Jorge M Finochietto. “Analysis of communication strategies for earth observation satellite constellations”. In: *IEEE Latin America Transactions* 14.6 (2016), pp. 2777–2782. DOI: 10.1109/TLA.2016.7555254. URL: https://doi.org/10.1109/TLA.2016.7555254.

[MFF18] Pablo G. Madoery, Juan A. Fraire, and Jorge M. Finochietto. “Congestion management techniques for disruption-tolerant satellite networks”. In: *Int. J. Satell. Commun. Netw.* 36.2 (2018), pp. 165–178. DOI: 10.1002/SAT.1210. URL: https://doi.org/10.1002/sat.1210.

[Mon+24] Sebastián I Montoya, Diego Maldonado, Juan A Fraire, and Sandra Céspedes. “On the Role of Delay Tolerant Networks and Contact Graph Routing in Direct-to-Satellite IoT”. In: *2024 IEEE 10th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2024, pp. 151–160. DOI: 10.1109/SMC-IT61443.2024.00024. URL: https://doi.org/10.1109/SMC-IT61443.2024.00024.

[NDF20] Andrea Nardin, Fabio Dovis, and Juan A. Fraire. “Empowering the Tracking Performance of LEO PNT by Means of Meta-Signals”. In: *8th IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2020, Venice, Italy, October 12-14, 2020*. IEEE, 2020, pp. 153–158. DOI: 10.1109/WISEE44079.2020.9262698. URL: https://doi.org/10.1109/WiSEE44079.2020.9262698.

[NFD22] Andrea Nardin, Juan A. Fraire, and Fabio Dovis. “Contact Plan Design for GNSS Constellations: A Case Study With Optical Intersatellite Links”. In: *IEEE Trans. Aerosp. Electron. Syst.* 58.3 (2022), pp. 1981–1995. DOI: 10.1109/TAES.2021.3135025. URL: https://doi.org/10.1109/TAES.2021.3135025.

[Olm+23] Matías Olmedo, Juan Andres Fraire, Renato Cherini, Fabio Patrone, and Mario Marchese. “Routing in Scalable Delay-Tolerant Space Networks with Graph Neural Networks”. In: *Proceedings of the 2023 International Conference on embedded Wireless Systems and Networks, EWSN 2023, Rende, Italy, September 25-27, 2023*. Ed. by Giancarlo Fortino, Valeria Loscrı`, Fabrizio Granelli, Tarek F. Abdelzaher, Antonella Guzzo, and Claudio Savaglio. ACM, 2023, pp. 211–218. DOI: 10.5555/3639940.3639967. URL: https://dl.acm.org/doi/10.5555/3639940.3639967.

[Ort+21] Raydel Ortigueira, Juan A. Fraire, Alex Becerra, Tomás Ferrer, and Sandra Céspedes. “RESS-IoT: A Scalable Energy-Efficient MAC Protocol for Direct-to-Satellite IoT”. In: *IEEE Access* 9 (2021), pp. 164440–164453. DOI: 10.

1109/ACCESS.2021.3134246. URL: https://doi.org/10.1109/ACCESS.2021.3134246.

[Ort+24] Raydel Ortigueira, Samuel Montejo-Sánchez, Santiago Henn, Juan A. Fraire, and Sandra Céspedes. "Satellite Visibility Prediction for Constrained Devices in Direct-to-Satellite IoT Systems". In: *IEEE Sensors Journal* (2024). DOI: 10.1109/JSEN.2024.3418728. URL: https://doi.org/10.1109/JSEN.2024.3418728.

[Ott+25] Camilla Ottaviani, Alessandro Compagnoni, Juan A. Fraire, Giacomo Verardo, Gabriel Maiolini Capez, Daniel Gaetano Riviello, Gregory Stock, Carla Fabiana Chiasserini, and Roberto Garello. "Advanced Routing Strategies for LEO and VLEO Constellations: Ensuring Polar Coverage". In: *14th Advanced Satellite Multimedia Systems Conference and the 20th Signal Processing for Space Communications Workshop, ASMS/SPSC 2025*. In Press. IEEE, 2025.

[Par+23] Pablo Ilabaca Parra, Samuel Montejo Sanchez, Juan A. Fraire, Richard Demo Souza, and Sandra Céspedes. "Network Size Estimation for Direct-to-Satellite IoT". In: *IEEE Internet Things J.* 10.7 (2023), pp. 6111–6125. DOI: 10.1109/JIOT.2022.3224678. URL: https://doi.org/10.1109/JIOT.2022.3224678.

[Pas+24] Hannaneh B. Pasandi, Juan A. Fraire, Sylvia Ratnasamy, and Hervé Rivano. "A Survey on Direct-to-Device Satellite Communications: Advances, Challenges, and Prospects". In: *Proceedings of the 2nd International Workshop on LEO Networking and Communication, LEO-NET 2024, Washington, DC, USA, November 18-22, 2024*. ACM, 2024, pp. 7–12. DOI: 10.1145/3697253.3697265. URL: https://doi.org/10.1145/3697253.3697265.

[Rav+18a] Fernando D. Raverta, Ramiro Demasi, Pablo G. Madoery, Juan A. Fraire, Jorge M. Finochietto, and Pedro R. D'Argenio. "A Markov Decision Process for Routing in Space DTNs with Uncertain Contact Plans". In: *6th IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2018, Huntsville, AL, USA, December 11-13, 2018*. IEEE, 2018, pp. 189–194. DOI: 10.1109/WISEE.2018.8637330. URL: https://doi.org/10.1109/WiSEE.2018.8637330.

[Rav+18b] Fernando D. Raverta, Pablo G. Madoery, Juan A. Fraire, and Jorge M. Finochietto. "Modeling Routing in Scheduled Delay Tolerant Networks Under Uncertainties". In: *Proceedings of the IV School on Systems and Networks, SSN 2018, Valdivia, Chile, October 29-31, 2018*. Ed. by Sandra Céspedes and Javier Bustos-Jiménez. Vol. 2178. CEUR Workshop Proceedings. CEUR-

WS.org, 2018, pp. 94–97. URL: https://ceur-ws.org/Vol-2178/SSN2018%5C_paper%5C_33.pdf.

[Rav+21] Fernando D. Raverta, Juan A. Fraire, Pablo G. Madoery, Ramiro A. Demasi, Jorge M. Finochietto, and Pedro R. D'Argenio. "Routing in Delay-Tolerant Networks under uncertain contact plans". In: *Ad Hoc Networks* 123 (2021), p. 102663. DOI: 10.1016/J.ADHOC.2021.102663. URL: https://doi.org/10.1016/j.adhoc.2021.102663.

[Roj+23] Camilo Rojas, Juan A Fraire, Fabio Patrone, Alberto Gotta, and Mario Marchese. "Continuous Time Emulation for Software-Defined Non-Terrestrial Edge Computing Networks". In: *European Wireless 2023; 28th European Wireless Conference*. VDE. 2023, pp. 230–235. URL: https://ieeexplore.ieee.org/abstract/document/10461445.

[Roj+24a] Camilo Rojas, Juan A. Fraire, Fabio Patrone, and Mario Marchese. "Advanced Constellation Emulation and Synthetic Datasets Generation for Non-Terrestrial Networks". In: *IEEE International Mediterranean Conference on Communications and Networking, MeditCom 2024, Madrid, Spain, July 8-11, 2024*. IEEE, 2024, pp. 37–43. DOI: 10.1109/MEDITCOM61057.2024.10621248. URL: https://doi.org/10.1109/MeditCom61057.2024.10621248.

[Roj+24b] Camilo Rojas, Juan A. Fraire, Fabio Patrone, and Mario Marchese. "Fuzzy Logic-Based Orchestration of Multi-Access Edge Computing in LEO Satellite Constellations". In: *IEEE International Conference on Communications Workshops, ICC 2024 Workshops, Denver, CO, USA, June 9-13, 2024*. IEEE, 2024, pp. 1685–1690. DOI: 10.1109/ICCWORKSHOPS59551.2024.10615305. URL: https://doi.org/10.1109/ICCWorkshops59551.2024.10615305.

[Roj+25] Camilo Rojas, Juan A. Fraire, Fabio Patrone, and Mario Marchese. "On the Latency Trade-off Between Space and Terrestrial Clouds in Non-Terrestrial Networks". In: *14th Advanced Satellite Multimedia Systems Conference and the 20th Signal Processing for Space Communications Workshop, ASMS/SPSC 2025*. In Press. IEEE, 2025.

[Sán+19] Daniel Sánchez, Pablo A. Ferreyra, Juan A. Fraire, Fabian Gomez, Raoul Velazco, and Dardo Vinas Viscardi. "A Wireless Embedded System for Measuring the Effects of Ionizing Radiations". In: *IEEE Latin American Test Symposium, LATS 2019, Santiago, Chile, March 11-13, 2019*. IEEE, 2019, pp. 1–6. DOI: 10.1109/LATW.2019.8704607. URL: https://doi.org/10.1109/LATW.2019.8704607.

[Sca+24] Germán L Scapin, Guido Alvarez, Nicolás Quiroga, Federico Collado, and Juan A Fraire. “Performance Assessment of LoRaWAN DtS-IoT Networks: Operation Modes and Server Placement”. In: *2024 IEEE Biennial Congress of Argentina (ARGENCON)*. IEEE. 2024, pp. 1–8. DOI: 10.1109/ARGENCON62399.2024.10735884. URL: https://doi.org/10.1109/ARGENCON62399.2024.10735884.

[SFH22a] Gregory Stock, Juan A. Fraire, and Holger Hermanns. “Distributed On-Demand Routing for LEO Mega-Constellations: A Starlink Case Study”. In: *11th Advanced Satellite Multimedia Systems Conference and the 17th Signal Processing for Space Communications Workshop, ASMS/SPSC 2022, Graz, Austria, September 6-8, 2022*. IEEE, 2022, pp. 1–8. DOI: 10.1109/ASMS/SPSC55670.2022.9914716. URL: https://doi.org/10.1109/ASMS/SPSC55670.2022.9914716.

[SFH22b] Gregory Stock, Juan A. Fraire, and Holger Hermanns. “Distributed On-Demand Routing for LEO Mega-Constellations: A Starlink Case Study”. In: *CoRR* abs/2208.02128 (2022). DOI: 10.48550/ARXIV.2208.02128. arXiv: 2208.02128. URL: https://doi.org/10.48550/arXiv.2208.02128.

[Sin+24] Arnau Singla Manau, Franco Criscola, David Canales, Anna M Calveras Augé, Juan Fraire, and Joan Adria Ruiz de Azua. “Integrating satellite constellation and mobile operations for non-terrestrial networks: preliminary results of dynamic scheduling”. In: *SPACOMM 2024, The Sixteenth International Conference on Advances in Satellite and Space Communications*. 2024, pp. 8–13. URL: http://hdl.handle.net/2117/415648.

[Sol+17] M. Solinas, Alexandre Coelho, Juan A. Fraire, Nacer-Eddine Zergainoh, Pablo A. Ferreyra, and Raoul Velazco. “Preliminary results of NETFI-2: An automatic method for fault injection on HDL-based designs”. In: *18th IEEE Latin American Test Symposium, LATS 2017, Bogotá, Colombia, March 13-15, 2017*. IEEE, 2017, pp. 1–4. DOI: 10.1109/LATW.2017.7906748. URL: https://doi.org/10.1109/LATW.2017.7906748.

[Sol+21] Germán J Soldano, Juan A Fraire, Jorge M Finochietto, and Rodrigo Quiroga. “COVID-19 mitigation by digital contact tracing and contact prevention (app-based social exposure warnings)”. In: *Scientific Reports* 11.1 (2021), p. 14421. DOI: 10.1038/s41598-021-93538-5. URL: https://doi.org/10.1038/s41598-021-93538-5.

[Sto+20] Gregory Stock, Juan A. Fraire, Tobias Mömke, Holger Hermanns, Fakhri Babayev, and Eduardo Cruz. “Managing Fleets of LEO Satellites: Nonlinear, Optimal, Efficient, Scalable, Usable, and Robust”. In: *IEEE Trans. Com-*

put. Aided Des. Integr. Circuits Syst. 39.11 (2020), pp. 3762–3773. DOI: 10.1109/TCAD.2020.3012751. URL: https://doi.org/10.1109/TCAD.2020.3012751.

[Sto+22] Gregory Stock, Juan A Fraire, Holger Hermanns, Eduardo Cruz, Alastair Isaacs, and Zhana Imbrosh. "On the automation, optimization, and in-orbit validation of intelligent satellite constellation operations". In: *SSC21-V-05, 2021 AIAA/USU Conference on Small Satellites, Logan, Utah, USA* (2022). URL: https://digitalcommons.usu.edu/smallsat/2021/all2021/168/.

[Sto+24] Gregory Stock, Juan A. Fraire, Santiago Henn, Holger Hermanns, and Andreas Schmidt. "A Stability-First Approach to Running TCP Over Starlink". In: *IEEE International Conference on Communications Workshops, ICC 2024 Workshops, Denver, CO, USA, June 9-13, 2024*. IEEE, 2024, pp. 1708–1713. DOI: 10.1109/ICCWORKSHOPS59551.2024.10615714. URL: https://doi.org/10.1109/ICCWorkshops59551.2024.10615714.

[VF20] Blas F Vega and Juan A Fraire. "Experimental Evaluation of On-Board Contact-Graph Routing Solutions for Future Nano-Satellite Constellations". In: *SSC20-WKVIII-08, 2020 AIAA/USU Conference on Small Satellites, Logan, Utah, USA* (2020). URL: https://digitalcommons.usu.edu/smallsat/2020/all2020/56/.

[VFH21] Kai Vogelgesang, Juan A. Fraire, and Holger Hermanns. "Uplink Transmission Probability Functions for LoRa-Based Direct-to-Satellite IoT: A Case Study". In: *IEEE Global Communications Conference, GLOBECOM 2021, Madrid, Spain, December 7-11, 2021*. IEEE, 2021, pp. 1–6. DOI: 10.1109/GLOBECOM46510.2021.9685152. URL: https://doi.org/10.1109/GLOBECOM46510.2021.9685152.

[Wal+24] Felix Walter, Marius Feldmann, Juan A Fraire, and Scott Burleigh. "The Architectural Refinement of μD3TN: Toward a Software-Defined DTN Protocol Stack". In: *2024 IEEE 10th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2024, pp. 161–170. DOI: 10.1109/SMC-IT61443.2024.00025. URL: https://doi.org/10.1109/SMC-IT61443.2024.00025.

[Wan+18] Yue Wang, Kanglian Zhao, Wenfeng Li, Juan A. Fraire, Zhili Sun, and Yuan Fang. "Performance Evaluation of QUIC with BBR in Satellite Internet". In: *6th IEEE International Conference on Wireless for Space and Extreme Environments, WiSEE 2018, Huntsville, AL, USA, December 11-13, 2018*. IEEE, 2018, pp. 195–199. DOI: 10.1109/WISEE.2018.8637347. URL: https://doi.org/10.1109/WiSEE.2018.8637347.

[Yan+20] Zhibo Yan, Juan Andres Fraire, Kanglian Zhao, Hongcheng Yan, Pablo G. Madoery, Wenfeng Li, and Hong Yang. “Distributed Contact Plan Design for GNSSs”. In: *IEEE Trans. Aerosp. Electron. Syst.* 56.1 (2020), pp. 660–672. DOI: 10.1109/TAES.2019.2917492. URL: https://doi.org/10.1109/TAES.2019.2917492.

[Yan+24a] Hongcheng Yan, Liang Qiao, Wei Wu, Juan A. Fraire, Dong Zhou, Luming Li, and Yong Xu. “Routing in future space-terrestrial integrated networks with SATNET-OSPF”. In: *Int. J. Satell. Commun. Netw.* 42.1 (2024), pp. 1–25. DOI: 10.1002/SAT.1495. URL: https://doi.org/10.1002/sat.1495.

[Yan+24b] Lei Yang, Juan A. Fraire, Kanglian Zhao, Ruhai Wang, Wenfeng Li, and Hong Yang. “Optimizing deep-space DTN congestion control via deep reinforcement learning”. In: *Comput. Networks* 255 (2024), p. 110865. DOI: 10.1016/J.COMNET.2024.110865. URL: https://doi.org/10.1016/j.comnet.2024.110865.

[Zho+19] Yan Zhou, Ruhai Wang, Lu Liu, Kanglian Zhao, Wenfeng Li, Xiangyu Lin, and Juan A Fraire. “Effect of Link Disruption on BP/LTP over a Relay-based Communication Architecture for International Space Station”. In: *2019 IEEE International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2019, pp. 95–99. DOI: 10.1109/SMC-IT.2019.00017. URL: https://doi.org/10.1109/SMC-IT.2019.00017.