ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Dottorato di Ricerca in
INGEGNERIA ELETTRONICA, TELECOMUNICAZIONI E
TECNOLOGIE DELL'INFORMAZIONE

Ciclo: XXX
Settore Concorsuale di Afferenza: 09/F2
Settore Scientifico Disciplinare: ING-INF/03

Interference Mitigation and Localization Based on Time-Frequency Analysis for Navigation Satellite Systems

Autore:
Giacomo POJANI

Supervisore:
Prof. Giovanni Emanuele CORAZZA
Coordinatore Dottorato:
Prof. Alessandro VANELLI-CORALLI

Dipartimento di Ingegneria dell'Energia Elettrica e dell'Informazione (DEI)
"Guglielmo Marconi"

2018

UNIVERSITY OF BOLOGNA

Abstract

School of Engineering and Architecture
Dipartimento di Ingegneria dell'Energia Elettrica e dell'Informazione (DEI)
"Guglielmo Marconi"

Doctor of Philosophy

Interference Mitigation and Localization Based on Time-Frequency Analysis for Navigation Satellite Systems

by Giacomo POJANI

Nowadays, the operation of global navigation satellite systems (GNSS) is imperative across a multitude of applications worldwide. The increasing reliance on accurate positioning and timing information has made more serious than ever the consequences of possible service outages in the satellite navigation systems. Among others, interference is regarded as the primary threat to their operation. Due the recent proliferation of portable interferers, notably jammers, it has now become common for GNSS receivers to endure simultaneous attacks from multiple sources of interference, which are likely spatially distributed and transmit different modulations.
To the best knowledge of the author, the present dissertation is the first publication to investigate the use of the S-transform (ST) to devise countermeasures to interference. The original contributions in this context are mainly:

- the formulation of a complexity-scalable ST implementable in real time as a bank of filters;
- a method for characterizing and localizing multiple in-car jammers through interference snapshots that are collected by separate receivers and analysed with a clever use of the ST;
- a preliminary assessment of novel methods for mitigating generic interference at the receiver end by means the ST and more computationally efficient variants of the transform.

These three topics are addressed in Chapters 2, 3, 4, respectively. The content of Chapter 2 is useful for any sort of application that requires the ST to process non-stationary signals of unknown characteristics. Besides GNSSs, the countermeasures to interference proposed are equivalently applicable to protect any direct-sequence spread spectrum (DS-SS) communication.
The research work described in this doctoral thesis or pursued on related topics during the three years of doctorate is either published or in preparation in [1–7].

Contents

List of Figures

List of Tables

List of Abbreviations

ACF	**A**uto-**C**orrelation **F**unction
AF	**A**mbiguity **F**unction
ADC	**A**nalog-to-**D**igitla **C**onverter
AGC	**A**utoomatic **G**ain **C**ontrol
AOA	**A**ngle **O**f **A**rrival
ARNSS	**A**eronautical **RNSS**
CAF	**C**ross-**AF**
CDMA	**C**ode **D**ivision **Multiple** **A**ccess
C/N_0	**C**arrier-to-**N**oise density power ratio
CRB	**C**ramér-**R**ao lower **B**ound
DME	**D**istance **M**easuring **E**quipmente
DBSCAN	**D**ensity-**B**ased **S**patial **C**lustering of **A**pplications with **N**oise
DoO	**D**evice **o**f **O**pportunity
DOST	**D**iscrete **O**rthonormal **ST**
DRSS	**D**ifferential **RSS**
DS-SS	**D**irect-**S**equence **S**pread **S**pectrum
EGNOS	**E**uropean **G**eostationary **N**avigation **O**verlay **S**ystem
EKF	**E**xtended **KF**
ENU	**E**ast-**N**orth-**U**p
ESA	**E**uropean **S**pace **A**gency
ESPRIT	**E**stimation of **S**ignal **P**aramters via **R**otational **I**nvariance **T**echniques
FAST	**F**requency-**A**daptive **ST**
FDOA	**F**requency **D**ifference **O**f **A**rrival
FI	**F**requency **I**nverse (ST)
FST	**F**ast **ST**
DFT	**D**iscrete **FT**
DS-SS	**D**irect-**S**equence **S**pread **S**pectrum
FBS	**F**ilter-**B**ank **S**ummation
FFT	**F**ast **FT**
FIR	**F**inite **I**mpulse **R**esponse
FT	**F**ourier **T**ransform
FWTM	**F**ull-**W**idth **T**enth-**M**aximum
GBAS	**G**round-**B**ased **A**ugmentation **S**ystem
GDOP	**G**eometrical **D**ilution **O**f **P**recision
GEMS	**G**NSS **E**nvironmental **M**onitoring **S**ystem
GFT	**G**enealized **F**ourier-family **T**ransform
GIDL	**G**eneralized **I**nterference **D**etection and **L**ocalization system
GNSS	**G**lobal **N**avigation **S**atellite **S**ystem
GSA	European **GNSS** **A**gency
HT	**H**ough **T**ransform
I	**I**n-phase
IDFT	**I**nverse **DFT**
IFFT	**I**nverse **FFT**

IF	**I**nstantaneous **F**requency
IIR	**I**nfinite **I**mpulse **R**esponse
IMM	**I**nteractive **M**ultiple **M**odel
IoT	**I**nternet **o**f **T**hings
ITU	**I**nternational **T**elecommunication **U**nion
FIM	**F**isher **I**nformation **M**atrix
JNR	**J**ammer to **N**oise power **R**atio
KF	**K**alman **F**ilter
LAAS	**L**ocal **A**rea **A**ugmentation **S**ystem
LNA	**L**ow **N**oise **A**mplifier
LS	**L**east **S**quares
MIMO	**M**ultiple **I**nput **M**ultiple **O**utput
ML	**M**aximum **L**ikelihood
MUSIC	**MU**ltiple **SI**gnal **C**lassification
NRMSE	**N**ormalized **RMSE**
OCXO	**O**ven **C**ontrolled **X**tal **O**scillator
OLA	**O**ver**L**ap and **A**dd
OS	**O**pen **S**ervice
PB	**P**ulse **B**lanker
PCRB	**P**osterior **CRB**
PF	**P**article **F**ilter
PDF	**P**robability **D**ensity **F**unction
APNT	**A**lternative **P**ositioning **N**avigation **T**iming
PRS	**P**ublic **R**egulated **S**ervice
PVT	**P**osition **V**elocity **T**ime
Q	**Q**uadrature
RF	**R**adio-**F**requency
RINEX	**R**eceiver **IN**dependent **EX**change format
RMSE	**R**oot **M**ean **S**quare **E**rror
RNSS	**R**adio **N**avigation **S**atellite **S**ervice
RSS	**R**eceived **S**ignal **S**trength
SBAS	**S**atellite-**B**ased **A**ugmentaion **S**ystem
SDR	**S**oftware-**D**efined **R**adio
SNR	**Signal** to **N**oise **R**atio
ST	**S**-**T**ransform
STFT	**S**hort-**T**ime **F**ourier **T**ransform
TACAN	**TAC**tical **A**ir **N**avigation system
TCXO	**T**emperature **C**ompensated **X**tal **O**scillator
TDOA	**T**ime **D**ifference **O**f **A**rrival
TF	**Time**-**F**requency
TI	**T**ime **I**nverse (ST)
TOA	**T**ime **O**f **A**rrival
TSST	**T**ime-**S**elective **ST**
UAV	**U**nmanned **A**erial **V**ehicle
UKF	**U**nscented **KF**
UTM	**U**niversal **T**ransverse **M**ercator coordinate system
WAAS	**W**ide **A**rea **A**ugmentation **S**ystem
WSN	**W**ireless **S**ensor **N**etwork
WT	**W**avelet **T**ransform
WVD	**W**igner-**V**ille **D**istribution

In memoria di Luciana Vasconi Gardi

Chapter 1

Introduction

Global navigation satellite systems (GNSS) deliver precise position and time information with global coverage. An ever-growing number of infrastructures and users, both civilian and military, relies on the operation of these systems. Applications relying on GNSSs include critical systems that span from the public to the private domains. Among the other, the accurate and worldwide positioning is an enabler of autonomous vehicles, which is a ground-breaking and safety-critical technology. More in general, driven by growing revenues, today and future market of navigation systems is globally expanding in several segments: location-based services, road, aviation, rail, maritime, agriculture, surveying, and synchronization. Besides positioning, timing is gaining more relevance, because modern networks in the telecommunications, energy, and finance sub-segments need a precise time reference. The increasing pervasiveness of GNSSs poses a significant concern about the vulnerabilities of these systems. Among them, the susceptibility to interference is arguably the major threat to their reliability. Awareness on this issue has grown since the publication of the VOLPE report in [8]. Existing regulations of the International Telecommunication Unions (ITU) prohibit the intentional broadcast of any unauthorized signal over the bands dedicated to aeronautical radio navigation services (ARNS) and radio navigation satellite services (RNSS). Therein lie the frequencies allocated to many navigation systems, such as the GPS, GLONASS, and Galielo, in chronological order. Despite the restricted access to these bandwidths, unauthorized transmissions are however occasionally observed by monitoring stations and proven by temporary and local outages. The daily occurrences of these interference events around infrastructures have steadily intensified in parallel with the number of GNSS users, arguably because of the growing privacy concern of the public toward the positioning service. To prevent these behaviours from causing damages to GNSS-enabled applications, countermeasures against interference have been extensively studied and documented over the last decades. Their sources may be split into two categories according to the purposes behind them. The first category encompasses unintentional and natural interference, which are caused by transmissions appearing at or near GNSS frequencies. Solar activity often causes powerful radiation bursts that can disturb large regions of the globe. Solar flares and other kinds of large-scale disturbances due to space weather increase the density of high-energy particle in the Earth's atmosphere, thence producing a sudden increment of background noise. When uninformed interference is man-made instead, it typically arises due to malfunctions of equipment designed to transmit in adjacent channels. Examples are the out-of-band harmonics seldom produced by AM/FM and amateur radios and TV repeaters. In this category falls also the intrinsic co-channel interference caused by the co-existence of other GNSSs as well as radar and augmentation systems. For instance, the impacts of the Distance Measuring Equipment (DME) and the Tactical Air Navigation (TACAN) aviation are studied in [9] for the L5/E5A bands. Another example discussed in [10] is the development of pseudo-satellites also known as pseudolites. The second category is that of sources of intentional interference, which is deliberately transmitted to disrupt or mislead the operation of GNSS receivers or, more rarely, satellites. Therein, spoofing and meaconing represent the more sophisticated and less common kinds of attack. Spoofers attempt to deceive a receiver by transmitting counterfeit GNSS signals slightly more powerful than the authentic ones. They forge a whole constellation of satellites with the aim of misleading the computation of the receiver position, velocity, and time (PVT) with consistency over time. The higher power causes the receiver to lock onto the false source rather than the true satellite. This kind of attack is effective against the open service (OS), because pseudo-random noise (PRN) codes are known and the navigation message is not encrypted. The use of encryption and classified codes of weekly length still prevents spoofers from threatening the public-regulated service (PRS), which is instead not inherently protected from meaconing. Meaconers carry out delay and re-transmit attacks with the aim of impairing the accuracy at the receiver end. They are transceivers that re-modulate the GNSS signals received to temper the inner delays and Doppler frequency shift. Similarly to spoofing, the counterfeit signals are then radiated to slightly overpower the authentic ones. Generally, the capability of elaborating fake navigation data and signals in real time is not trivial and can target one victim at the time. Hence, although the availability of entry-level software-defined radios (SDR) is changing the game as proven in [11], devices equipped with spoofing or meaconing functionality are not yet within the reach of the public. On the contrary, jamming attempts simply consist in brute-force broadcast of powerful interference, which can be generated with basic hardware and no specific software. Their goal is to cause the denial of both OS and PRS over a certain area, which can be accomplished regardless of the number of receivers present. Interfering with a satellite in space requires the radiation of significant amount of power through large-size and highly-directive antennas, which are easy to spot and hard to get. Therefore, the most prevalent and thus dangerous attacks target terrestrial receivers. Indeed, since GNSS signals have extremely low power levels (i.e., about -160 dBW) once they arrive on the Earth's surface, they are likely overpowered by any source of interference in the surroundings. Particularly, the deliberate transmission of signals at or near GNSS frequencies without the desire to cause harm to the system itself or to manifold receivers is usually referred to as uninformed. An interesting example of this king are so-called Personal Privacy Devices (PPD) in [12]. Although their usage is illegal, they are gaining popularity among the public due to the privacy concerns related to the localization services and their advertised range of effectiveness is often underestimated: tens of meters instead of hundreds. This misunderstanding is likely to happen because most of the jammers on the market are designed to interfere also with wireless communications (i.e., GSM, UMTS, LTE, WiFi, etc.) and other tracking systems (e.g., LoJack). These functionalities go far beyond the protection of user's own privacy. In order to harm all these systems at the same time, the devices commercialized are actually arrays of jammers, each one provided with a separate antenna suitable to jam a certain band. A sample of these devices are shown in Fig. 1.1. General-purpose jammers as such are illegal, but they may be purchased online and shipped from countries with loose regulations. They are typically able of radiating 1 W per channel, the sub-band of which is tens of megahertz wide and is spanned in a matter of few microseconds. In the last decade, low-cost and hand-held versions of these devices became more and more widespread on board of civilian vehicles. This trend has been demonstrated by fieldworks and accidents. On the one side, stations used for field investigations and for augmentation systems have been increasingly observing jamming attempts across the road network, as documented in [13, 14]. On the other side, outages of the GNSS operation have also being experienced on a daily basis by critical infrastructures located near heavy-traffic highways, such as airports and harbors. In the following sections, some relevant reports are summarized.

FIGURE 1.1: Sample of jammers for sale online in ascending order of complexity, power and price (jammers.it).

1.0.1 Relevant Interference Events of Public Record

- In late 2009, the ground-base augmentation system (GBAS) of the Newark Liberty International Airport, United States, was suffering from brief daily breaks due the uninformed usage of single in-car jammer. This system enhances the integrity and accuracy of the positions estimated by the GPS constellation in order to support the approach, landing, and departure of aircrafts and to avoid congestions in the air traffic. The illegal transmission was originated on a truck often passing onto the highway nearby, as told in [15]. The driver intended to block the on-board GPS-based tracking device in order to prevent it from communicating to the fleet manager his position and speed. However, this trick turned out to accidentally harm also other GPS signals in the area, including those used by air navigation systems. At that time, the device was regularly sold as a cheap PPD at the cost of about 30 USD. In the GBAS, indeed, all reference stations have antennas typically located within 100-200 meters from each other. Therefore, after this episode, a measure for mitigating but not eliminating the damage due to the daily jamming attempts was to modify the sites and radiation pattern of the antennas. Also satellite-based augmentation systems (SBAS) have been experiencing the same issue, like the wide area augmentation system (WAAS) and the European geostationary navigation overlay system (EGNOS). Nevertheless they are more robust to the harming effects of interference, simply because they can afford the temporary loss of individual reference stations that are many and widely spread geographically.

- In January 2007, GPS services were disrupted throughout San Diego, California, United States. The naval medical center emergency pagers stopped working, the harbor trace-management system used for guiding boats failed, the airport control had to resort to non-GPS backup systems, mobile-phones users found themselves out of coverage, and bank customers were refused while trying to withdraw money. The origin of this inconvenience was attributed to two military navy ships that were conducting a training exercise in the bay and jammed radio signals by accident.

- In April 2012, in Kent, England, the police arrested members of a criminal gang responsible for the theft of some 150 high-value vehicles, using jammers to disable the in-car tracking systems cars.
- In April 2013, North Korea is believed to have broadcast jamming signals that neutralized GPS navigators on at least 250 flights in South Korea. The transmit power radiated was estimated to be about 50 W. The interference was enough powerful to cause outages in the mobile network of Seoul, which is tens of kilometers away from the border.
- In November 2013, in Australia, a Melbourne newspaper reported that more than 100 cabs in the city were suspected of using GPS jammers in order to fool the fleet management into giving them jobs, even though they were not in the area. The devices were discovered as they were also obscuring the location of nearby GPS-equipped emergency vehicles, like police cars, ambulances, and fire trucks.

(From magazine articles in Inside GNSS and GPS World)

1.0.2 Efforts against GNSS Interference in Europe

Completed in March 2012 and run until October 2013, the DETECTOR project was supported by European Commission funding through the European GNSS Agency (GSA) and motivated by the need to "fingerprint" interference threats to GNSS-based road applications, with focus on jamming. More information are provided within the deliverable in [16] or may be found at gnss-detector.eu. The consortium carrying out the project included the University of Bologna for research of interference countermeasures, and NSL as technology developer. An initial investigation of interference attempts was carried out using datasets recorded from a network of over 100 roadside monitoring stations in the United Kingdom and Ireland. In these sites, many potential interference events in terms of temporary and unexpected large observations errors were detected by comparing the estimated position against the known reference coordinates as well as by recognizing drops in the power levels received from multiple satellite at the same time. For instance, the preliminary analysis of data collected over two days at the site close to the London identified more than 20 separate attacks with a variety of stationary and non-stationary signatures, hence providing preliminary evidence of the spread use of jammers in the road network. The ultimate task of these probes is to continuously monitor the spectrum and build a database of all interference events on a back-end server. The adoption of software receivers has made possible to autonomously characterize jamming waveforms at the digital sample level in order differentiate among unintentional and different kinds of intentional interference sources directly at the target sites. An evolution of the working prototype of the probe is currently commercialized by Spirent and NSL: the GSS200D depicted in Fig. 1.2. This device is meant to report and analyze the interference activity at site of interest over multiple GNSS bands.

FIGURE 1.2: GSS200D interference detector by Spirent (spirent.com/Products/GSS200D-Detector).

Running between 2008 and 2014, the GAARDIAN earlier and the SENTINEL later were research programs part-funded by the United Kingdom government within a collaboration of universities and institutions led by Chronos Technology. The objective was investigating the vulnerability to interference of a number of mission-critical and safety-critical services relying on GNSS, examining the growing use of low-cost jammers, and creating a nationwide system to provide intelligence for law enforcement actions against their usage. In other words, the aim was prove the technology is ready to tackle jamming attempts but also to prove the threat, thereby validating the market as well as the technical countermeasures. For this reason, a mesh network of probes known as Signal Sentry 1000 was deployed to detect, measure, and geolocate any interference onto the GNSS bands, discriminating between space weather effects and intentional attacks. The reports produced by this network demonstrated the worsening of the jamming phenomenon: some probes detected 5 to 10 events per day with increasing trend, while more and more powerful devices were recorded as their average ranges considerably grew over the measurement campaign. Moreover, the project undertook a survey to determine the extent of the market of GPS jammers on sale over the Internet. Over 50 websites were found actively selling these illegal devices, which were then purchased to assess their specifications and to test GPS receivers. The preliminary conclusions disclosed in [17] highlight that any low-cost jammer (i.e. on average 50 USD) within 200 meters could compromise any GNSS-based application in absence of countermeasures. Indeed, whereas GPS was the only system under threat at the time of the Volpe report in [8], by 2014 all frequencies allocated to GNSSs became under attack. Since 2013, the same partnership between industry and academia worked on a third project for automatic jamming recognition. The latest news in May 2015 regard the demonstration of a proof-of-concept unit, which triggers a camera whenever an in-car jammer is detected for the vehicle passing in front of the lens, and transfers a photo to the central server accessible to police officers.

Funded under Horizon 2020 by the European GSA, STRIKE3 is an ongoing project for the development of international standards in the area of GNSS threat monitoring and reporting, and receiver testing. One the one hand, standardized formats are necessary for the operation of a worldwide GNSS interference monitoring network, which captures and stores all the events recorded by using different equipment and sensors within a common shared database. After capturing, the data of interference events appearing all over the world can be used to test GNSS receivers and enhance their robustness. The ultimate goal is to provide the common technical ground for police forces, highways authorities, toll operators, ports authorities and governmental organizations to create an international task-force in order to discourage the criminal use of jammers and contain the damages. More information are available at aic-aachen.org/strike3.

1.0.3 Focus of the Present Dissertation

Jamming differs from meaconing and spoofing because of the simplicity and the effectiveness in conducting these kinds of attacks. This fact makes jammers arguably the main threat to systems vulnerable to interference, such as GNSSs, and it also explains their massive availability on the black market at low retail prices. Nowadays, the extent of the proliferation of portable and so movable jammers has increased the danger of undergoing attacks from many sources of interference at the same time. Since the powerful interference generated by jamming is by nature exposed to the victim receivers, a plethora of detection mechanisms have been investigated in the literature and implemented into several commercial receivers for various technologies, from cellular to satellite communications. In other words, flagging a jamming attempt is a functionality well-consolidated from the theoretical and practical standpoints in both academia and industry. Instead, taking active measures to neutralize the detected attackers remains an open issue. Therefore, in the context of countermeasures to interference, active research efforts have now incremented towards the provision of cooperative monitoring systems capable of localizing multiple jammers as well as add-ons to receivers for mitigating simultaneous jamming waveforms. Both these problems are the subject of the present doctoral thesis, which explores signal processing solutions that could be implemented without expensive and bulky antenna systems. As far as GNSSs are concerned, it is worth mentioning that research lines parallel to the topics of this dissertation focus on backup systems inherently robust to interference and outages. Alternative positioning, navigation, and timing (APNT) systems and multi-sensor fusion are two prominent examples in this area.

Each chapter is introduced with the relative state of the art, stressing the limitations of the existing methods. Chapter 2 introduces the reader to consolidated and innovative mathematical notions necessary to understand the countermeasures to interference presented afterwards. Chapters 3 and 4 deal with the localization and the mitigation of multiple jammer, respectively. The intent is to demonstrate the promise of time-frequency (TF) analysis in neutralizing the presence of many civil jammers, which threat the operation of direct-sequence spread spectrum (DS-SS) and GNSS receivers.

Chapter 2

Time-Frequency Analysis

The scientific literature includes a large variety of mathematical tools dedicated to the analysis of signals in a transformed time-frequency (TF) domain. The ST provides a linear TF representation that identifies both the amplitude and the phase of the components that are otherwise blended in the Fourier spectrum. As opposed to those provided with other transform, the representations at the output of the ST can have *consistent* TF resolution regardless of the inputs, thanks to the use of multiple analysis windows. Such a capability, however, comes at an excessive price in terms of *complexity*, which has limited the potential utility of the ST to a few applications mainly in physics and biomedicine. Generally speaking, the ST can be used for post-processing relatively *short* amount of data *offline*. As the dataset size increases, though, the amount of processing power necessary to maintain the same computational time becomes rapidly prohibitive. For this reason, scientists and engineers are usually forced to abandon the ST in favour of the short-time Fourier transform (STFT), which relaxes the computational requirements but is intrinsically limited in terms of TF resolution due to the adoption of a single window. In particular, the STFT is the standard tool for *real-time* applications, including methods for interference mitigation and localization.

In order to extend the utility of the ST, the author of the dissertation contributed to the development of a novel mathematical framework for the forward and inverse computations of a ST that is scalable in terms of complexity and, consequently, potentially suitable to real-time execution. To this regard, we have formulated a new and generalized TF sampling scheme with scalable redundancy both in time and frequency. Contrary to precedent non-redundant (i.e., one-to-one) "fast" sampling schemes for the ST, ours provides an arbitrary degree of flexibility to trade computational efficiency for accuracy, or vice versa. As we prove with a case study, this freedom is fundamental when the TF representation is filtered to modify the signal spectral content over time. In the following sections, a formulation of the ST as a bank of filters is firstly presented by analogy with the STFT. Secondly, a generalized sampling for re-scaling the amount of redundancy embedded into the ST is devised within the novel framework in order to reduce the complexity of the filter-bank architecture as needed. The ultimate goal is to design a practical and flexible tool, which is able of analyzing, filtering, and synthesizing non-stationary signal and suitable to real-time processing.

In anticipation of the next sections, the following list is an overview of the original contributions that may be found in more detail throughout the chapter:

- a comprehensive analogy between ST and STFT, which made possible the migration of some concepts from one transform to the other;

- the convenient interpretation of the discrete ST as a bank of parallel digital processing chains, each of which has a down-converter and a lowpass Gaussian filter having bandwidth progressive with respect to the modulation frequency;

- a discussion about the potential to undersample the redundant TF representation of the ST;
- a discussion about the impact of the side effects of the progressive size of the windows of the ST;
- the explicit formulation of the constraints that implicitly underlie the inversions of the ST, which are then useful to scale down the complexity of the transform without losing information in the process;
- the addition of a simple correction in the discrete ST definition in order to grant the previously missing "neutrality" of the transform with respect to the time and frequency versions of the forward computation;
- direct and indirect procedures for reconstructing a generic waveform through the instantaneous frequency estimation based on the ST, as an alternative to the inversion for mono-component signals;
- the formulation of the complexity-scalable architecture of the ST based on a generalized TF sampling scheme, which improves on the existing non-redundant schemes in terms of flexibility and which was the insight driving of the entire work done on the topic;
- the formulation of an "octave reverse" that specifically inverts the minimally-sampled ST accounting for a phase correction, which was missing in precedent papers;
- the study of a case of an exemplary TF filter, where the flexibility of our scheme turns out to be crucial in achieving a desired compromise between the computational efficiency of the input analysis and the accuracy of the output synthesized from the filtered TF representation.

All in all, the concepts explored in this chapter are important to understand the novelty and the effectiveness of methods for interference mitigation and localization proposed in Chapter 3 and 4, where the use of the ST is made for the first time to overcome the issues of the other approaches in the state-of-the-art (e.g., based on the STFT). Most of all, any practical hardware or software implementation of these methods inside GNSS receivers or (cloud-computing) servers could lie on the foundations of the complexity-scalable architecture formulated here, by trading some accuracy to execute any ST-enabled algorithm in a timely manner. Nearly real-time processing through the ST is now more easily within the reach, depending on the computational resources at disposal and the dimensions of the filter-bank architecture. Within the mathematical framework described, the scalability of the ST could be used to the benefit of a number of scientific applications of any field, which were previously obliged to adopt the STFT for analysing, filtering, or synthesizing non-stationary signals in a reasonable time.

2.1 State of the Art

The spectral analysis of signals is a crucial task in countless scientific applications. The conventional Fourier transform (FT) decomposes a stationary signal into its individual frequency components, but it is unsuitable to represent dynamic spectral contents. The need for properly analyzing non-stationary signals has motivated the development of transforms over two-dimensional domains. Defined as such, they can catch both the temporal and spectral variations of the signal components. This field is referred to as TF analysis. Given the extension of the topic, hereinafter we briefly recall only on the most conventional alternatives to the ST. For a comprehensive review of the literature, the reader may refer to [18, 19] and the references cited therein. Among the manifold representations studied throughout the last decades, the STFT is certainly one of the most popular ones, both in theory and practice. Together with the ST and others variants of the FT, it may be related to the general Fourier-family transform (GFT) defined in [20]. This family encompasses transforms that are linear and constructed by using the products of windows and sinusoidal kernels as basis functions. The simplest member is obviously the FT, where the window is just a unitary coefficient. On top of this transform, various strategies of windowing allow us to determine the timings behind the amplitudes and the phases of dynamic spectra, but they also impose a trade-off between temporal and spectral resolutions in the TF representation. As explained in [21], such a trade-off is an inevitable consequence of the uncertainty principle of Heisenberg, which states the physical impossibility of simultaneously observing quantities that are complementary by nature, like time and frequency, with the same degree of precision. For instance, the STFT analyzes a signal by sliding a window to capture local changes in the spectrum with short terms of fixed resolution, as explained in more detail in [22]. Therefore, the resultant representation is conditioned by the choice of the single analysis window, which sets a constant trade-off between the temporal resolution at the slow frequencies (near the origin) and the spectral resolution at the fast frequencies (closer to the filter cut-off frequency). The ST may be regarded as the extension of the STFT to multi-resolution analysis, because it employs frequency-dependent windows. In the original formulation of [23], these windows are Gaussian functions with standard deviation directly proportional to the frequency. Thanks to them, the ST improves on the STFT by achieving a progressive trade-off in the TF resolution. Although the Gaussianity has the convenient property of minimizing the time-bandwidth product, which is lower bounded by the uncertainty principle, there are actually many functions that may be adopted for windowing. Their various impacts on the identification of harmonics is studied in [24]. Alternatively to the GFT, other popular TF representations are based on the wavelet transform (WT) of [25] and the Cohen's class of bilinear TF distributions, of which the Wigner distribution in [26] is a well-known member. In principle, the WT decomposes a signal into a set of self-similar series of dilations and translations of a mother wavelet. Similarly to the ST, it can locate the amplitude of the global spectrum with progressive resolution. However, it cannot provide the *globally-referenced* phase (i.e., referred to the time instant zero) that is instead retained by the GFT. In this context, a comparison between the ST and the WT is made in [23]. As far as the extensive Cohen's class of TF distributions reviewed in [27] is concerned, the instantaneous auto-correlation function (ACF) of the signal is used to analyze the respective time-varying power spectra in the TF plane. Because of the quadratic nature of these distributions, the dynamic spectra are represented with superior resolution that is not subject to the aforementioned trade-off. Nevertheless, the representations suffer from the presence of cross terms, which are nothing but artifacts produced by the interactions between the TF components of the signal that enter twice the ACF. The presence of these terms could obscure or mislead the interpretation of the signal characteristics. For this reason, bilinear distributions are not good candidates for the analysis of multi-component signals, if no a-priori knowledge is available about the individual incoming waveforms. A comparison between TF representations is provided in Fig. 2.1, where the ST clearly benefits from the progressive trade-off with good average resolution and no cross terms.

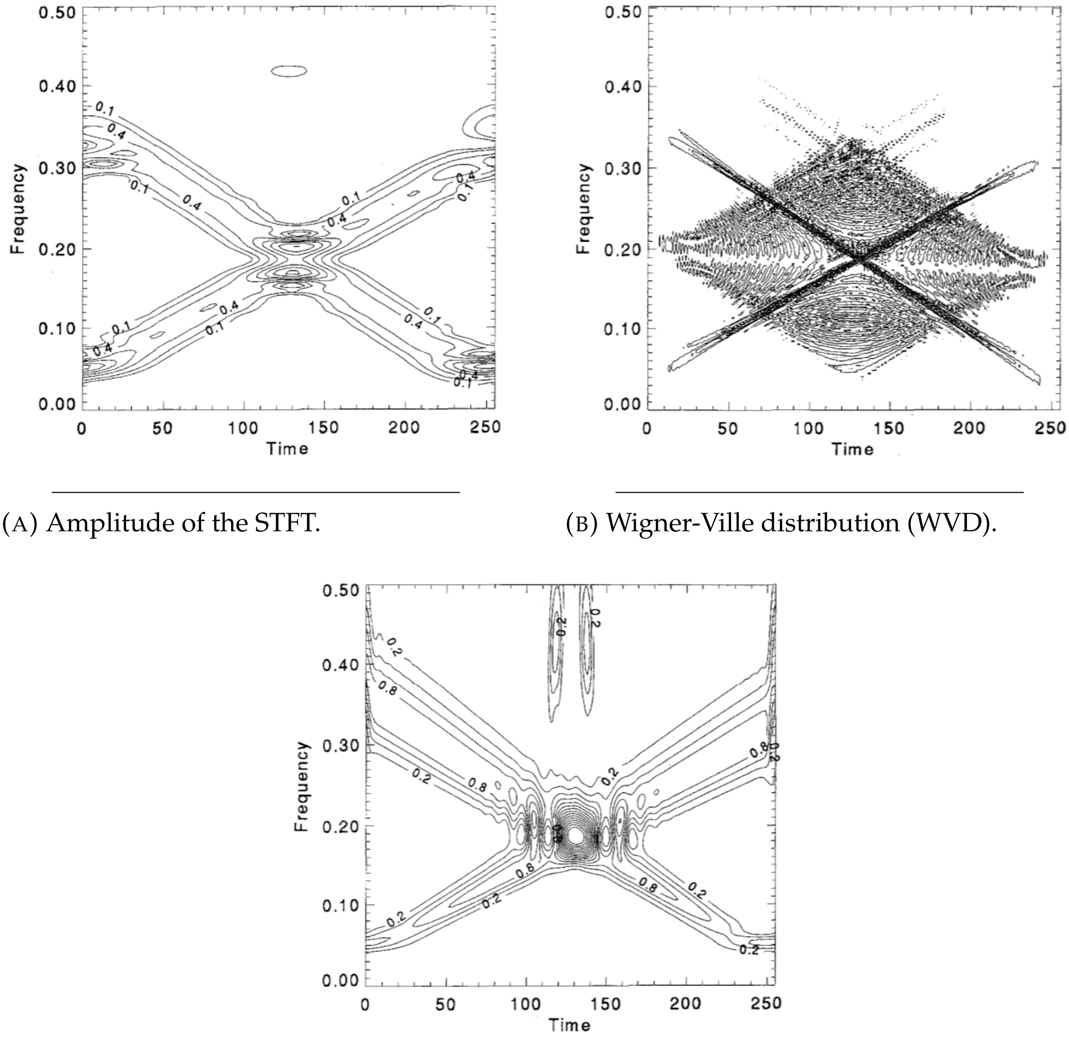

(A) Amplitude of the STFT.

(B) Wigner-Ville distribution (WVD).

(C) Amplitude of the ST.

FIGURE 2.1: Example in [23] for the TF representation of the superimposition of two crossing linear chirps and two high-frequency bursts.

Among the mathematical tools in the realm of TF analysis, the ST uniquely combines **linearity** and **multi-resolution** properties to locate both the amplitude and the globally-referenced phase of the signal in the TF domain. Although these appealing features have shown promise in various fields, such as in [3, 28–31], just to name a few, some issues related to the ST are still open for research. The main issue is the demanding requirements in terms of processing power and storage space, which hinder the feasibility of the timely computation of this TF representation with little computational resources. This aspect is especially critical when latency constraints should be fulfilled. Consequently, the complexity has restrained the field of application of the the original transform to the post-processing of relatively short datasets, while real-time applications usually resort to standard versions of the STFT or to the WT. Motivated by the potential utility of the ST in emerging applications, several attempts have been proposed to alleviate its computational and storage burdens. In this context, a discrete orthonormal ST (DOST) and the analogous fast ST (FST) are proposed in [32] and [20], respectively, to address the issue of complexity. In particular, the FST makes use of Gaussian windows that are separate and contiguous in frequency according to a dyadic scale and truncated in time with durations sized accordingly. The intent of such a minimal sampling scheme is to design a one-to-one representation free from redundant information. Both the DOST and the FST accomplish this task and thus minimize the complexity of this transform. Although they are both coherent with the progressive TF resolution underlying the intrinsic TF trade-off, they might also hide meaningful characteristics of the signal analyzed that could be exposed with little more redundancy. Another limitation is that one-to-one TF representations are not suitable to be filtered with accuracy and without introducing excessive distortion. In other words, the TF filters result poorly selective and the time-domain signals synthesized out of them are extremely distorted. This downside of the DOST/FST was likewise anticipated for the non-redundant STFT in [33] and we will later evaluate its effect with an example application. Because of this downside, other signal-specific strategies may be preferable when the complexity has to be reduced. For example in order to characterize the disturbances affecting power lines, a simple one-dimensional sampling scheme is used in [34] to compute the ST only for the fundamental frequency and its harmonics. Equivalently, a novel method for interference rejection is described in [3] and Chapter 4: the analysis based on the ST is restricted to the sole signal components that concentrate most of the incoming energy spectral density. Both these works aim at speeding up the computation of the transform by selectively processing the portions of the TF representation with noticeable energy. However, by doing so, they save computations by an extent that depends on the signal itself and cannot be bounded *a priori*. In other words, the complexity reduction is not guaranteed. Besides, a second issue lies in the invertibility of the discrete ST. In more detail, two approaches have been proposed in the literature to invert the continuous ST; one is referred to as the frequency inverse (FI) in [23] and is exact, whereas the other one is known as the time inverse (TI) in [35] and provides an approximation, which is proven in [36]. Once applied to finite-duration and discrete-time samples, the previous inverses are not neutral. In fact, they produce artifacts and their efficiency change, depending on how the forward transform is implemented in the first place. This troublesome aspect is discussed also in [37] as well as in the next sections.

To the best knowledge of the authors of [6], the derivation of a *flexible* architecture of the ST in terms of complexity is an original contribution to the state of the art and it is particularly useful for real-time applications. The conclusions of this work led to the draft of a journal paper that is currently submitted for publication. Therein, our goal is to retain the unique and advantageous features of this transform with controllable computational efficiency. First of all, we devise the ST as a *bank of filters* by porting the knowledge about the STFT to this more general context. The extension is straightforward and turns out to be a necessary step in deepen the understanding about other relevant notions of the ST, such as the neutrality with respect to the forward computation. This perspective also allows us to define the necessary and sufficient conditions for the exact signal recovery, either through the FI or the TI. Furthermore, it lays down the foundations for real-time implementation of the transform. The primary question in this respect is how to efficiently sample the TF plane underlying the original and fully-redundant ST. The answer is a **complexity-scalable** transform that produces the one-to-one dyadic scheme as a special variants. The amount of redundancy embedded into the TF representation is re-sizable, so that the analysis is performed with an arbitrary trade-off between complexity and accuracy. The secondary question is how to retrieve back the signal from the sampled representation without losses of information and as accurately as possible. Reconstructing a sequence of samples is possible either by inverting the transform or by inferring the signal amplitude and phase from its instantaneous frequency (IF), which is estimated from the magnitude peaks in the ST. Finally, the combination of the scalable filter-bank architecture with a reconstruction based on inversion/estimation provides a complete process to analyze, modify, and synthesize non-stationary signals. In particular, the flexible sampling scheme is validated in case study, which evaluates the effectiveness of TF filtering against the corresponding computational complexity.

2.2 The S-Transform as a Bank of Filters

For the sake of practicality, we shall restrict ourselves to discrete time series. In this context, a filter-bank interpretation of the discrete ST is derived from that of the STFT, which is briefly recalled from [38]. The presented architecture of the ST has the potential to target tight delay constraints, while overcoming the limitations of the STFT.

2.2.1 Short-Time Fourier Transform

Let us denote as $x[n]$ a complex-valued time series sampled at a rate f_x and assume an even number N of samples. The discrete STFT of $x[n]$ is defined by

$$\mathrm{STFT}_x\left(\frac{n}{f_x}, \frac{p}{F_x}\right) = \mathrm{STFT}_x[n,p] = \sum_{m=0}^{N-1} x[m]w[n-m]\mathrm{e}^{-j\frac{2\pi}{N}pm} \tag{2.1}$$

for $n = 0, 1, ..., N-1$ and $p = -N/2, 1, ..., N/2-1$, where n and p denote the time and the frequency indices, respectively, f_x and F_x are the respective sampling rates in the time and frequency domains, and $w[n]$ is the analysis window. Eq. 2.1 defines a square matrix of order N, the elements of which are TF components sampled every time period $\Delta t = 1/f_x$ and frequency step $\Delta f = 1/F_x = f_x/N$. The window is a discrete and non-zero function of length N_w, and assumed to be real and symmetric. A *short-time section* of $x[n]$ at the time index n_0 is defined as

$$x_{n_0}[n] = x[n]w[n_0 - n] \tag{2.2}$$

that is the product of $x[n]$ and the window time-reversed, and time-shifted by n_0 samples. By fixing the time index to n_0, the summation in Eq. 2.1 may be expressed as the discrete FT (DFT) of $x_{n_0}[n]$:

$$\mathrm{STFT}_x[n_0, p] = \mathrm{DFT}_n\{x_{n_0}[n]\} = \sum_{m=0}^{N-1} x_{n_0}[m]\mathrm{e}^{-j\frac{2\pi}{N}pm}. \tag{2.3}$$

Therefore, the transform is computed by performing the DFT for a set of N overlapping short-time sections, which are obtained by sliding the window sample-by-sample over the signal analyzed. Equivalently, by fixing the frequency index to p_0, the summation in Eq. 2.1 may be expressed as the circular convolution of $w[n]$ and the signal that is heterodyned by frequency-shifting $x[n]$ from the digital frequency

p_0/N to zero:

$$\mathrm{STFT}_x[n,p_0] = \left(x[n]\mathrm{e}^{-j\frac{2\pi}{N}p_0 n}\right) \circledast w[n]. \tag{2.4}$$

The complex exponential modulator performs a frequency down-conversion, whereas the window acts as the impulse response of a prototype lowpass filter. The combination of these two operations is the same as bandpass filtering the input signal. Hence, the whole STFT is carried out through a bank of N parallel channels splitting the Fourier analysis into overlapping sub-bands, which are equally spaced about the unit circle.

Although the DFT and the filtering formulations of the STFT are theoretically equivalent, they practically differ in their implementations. For instance, a fast FT (FFT) block in Eq. 2.3 might replace the need for a convolver, which would be instead realized as a finite impulse response (FIR) filter in Eq. 2.4. These two implementations correspond to two different choices in terms of latency and complexity. On the one hand, the FFT block has to buffer the incoming samples, before efficiently processing them. The additional latency between these block-by-block computations might prevent real-time applications from being feasible for large blocks of samples. On the other hand, the FIR filter outputs one sample per unit time, as soon as the input signal propagates through the tapped delay line, hence after a certain transient response. However, if a large number of taps is required, the poor computational efficiency of a direct-form FIR filter might be prohibitive. Hybrid strategies trade computation time for lower requirements, such as that in [39]. On the contrary, the architecture proposed in what follows does not compromise on latency: we analyze the signal with high-speed parallelism through a bank of digital filters, the feasibility of which is so a matter of complexity.

2.2.2 S-Transform

Following the formalism in [20], the discrete GFT of $x[n]$ may be written as

$$\mathrm{GFT}_x[n,p] = \sum_{m=0}^{N-1} x[m]w[n-m,v]\mathrm{e}^{-j\frac{2\pi}{N}pm} \tag{2.5}$$

where $w[n,v]$ consists in a set of N windows, which are functions with respect to the time index n and the generic variable v. If the windows are frequency-dependent (i.e., $v = p$), Eq. 2.5 coincides with that of the generalized ST in [40], whih includes the FT and the STFT. Let us consider the ST as originally presented in [23]. This transform employs Gaussian windows given by

$$w[n,p] = \frac{|p|}{N\sqrt{2\pi}}\mathrm{e}^{-\frac{1}{2}\left(\frac{np}{N}\right)^2}, \qquad p \neq 0 \tag{2.6}$$

that have the same size as the input time series (i.e., $N_w = N$) and standard deviation inversely proportional to the frequency (i.e, $\sigma_p = N/|p|$). Their similarity with the wavelets come from the dilations/contractions of a Gaussian function. In Eq. 2.6 a special definition is evidently necessary for the zero-frequency window:

$$w[n,0] = \frac{1}{N} \tag{2.7}$$

which acts on $x[n]$ as a moving average filter. In comparison to any window used in the STFT, the set of windows defined by Eqs. 2.6 and 2.7 enhances the spectral resolution at low frequencies, by covering more oscillation periods the sinusoidal

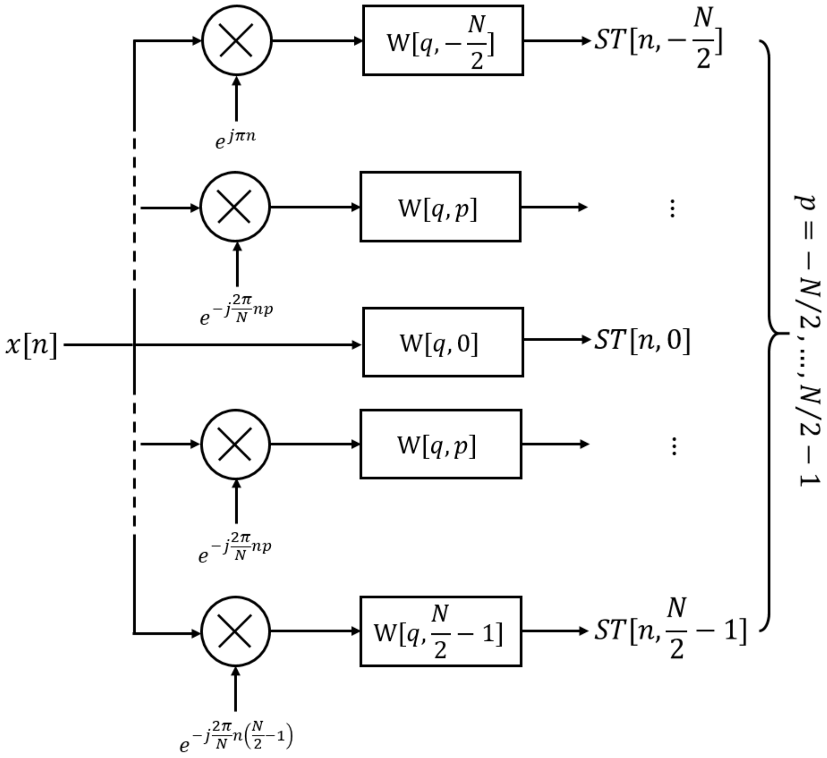

FIGURE 2.2: Forward ST computed through a bank of down-converters and lowpass Gaussian filters.

basis functions with wider lobes, as well as the temporal resolution at high frequencies, by limiting the time spread into shorter lobes. Consequently, like the WT, the ST achieves a progressive TF resolution trade-off, from which the resultant TF representation benefits in terms of consistency: multiple and simultaneous waveforms with different TF characteristics are visible thanks to the progressive resolution. On the contrary, the STFT performs a fixed-resolution analysis with a single window that is suitable only to signals of specific characteristics. More specifically, in [41], it is demonstrated that the window can be optimized to represent the linear chirps of a certain rate. More generally, the accuracy of the STFT is highly sensitive to the compatibility of the window with the signal under analysis, as argued in [42]. The multi-resolution extension of this transform is obtained by elaborating on Eqs. 2.3) and 2.4. The ST of $x[t]$ with constant time index n_0 is the summation of the signal multiplied with the time-shifted window and the basis functions:

$$\mathrm{ST}_x[n_0,p] = \sum_{m=0}^{N-1} x[m]w[n_0-m,p]\mathrm{e}^{-j\frac{2\pi}{N}pm} \tag{2.8}$$

which cannot be expressed through the DFT, due to the dependence of $w[n,p]$ on frequency, as opposed to Eq. 2.3. The ST of $x[n]$ at a specific frequency p_0 is known as a *voice* that is the convolution of $w[n,p_0]$ and the signal down-converted by the digital frequency p_0/N:

$$\mathrm{ST}_x[n,p_0] = \left(x[n]\mathrm{e}^{-j\frac{2\pi}{N}p_0 n}\right) \circledast w[n,p_0] \tag{2.9}$$

similarly to Eq. 2.4. Therefore, the discrete ST may be formulated as a bank of digital filters, in which each of the N uniformly-spaced channels computes a voice that isolates one frequency component of $x[n]$. The fact that the lowpass filters are not identical and have bandwidths increasing with $|p|$ distinguishes this interpretation from that of the STFT. Accordingly, the voice in Eq. 2.9 is one of the outputs of the aforementioned filter-bank, which is represented in Fig. 2.2. The direct calculation of all voices entails an overall complexity on the order of $O(N^3)$ operations and $O(N^2)$ memory elements according to [20]. Such a load can be handled by N FIR filters of length N each. Otherwise, a formulation equivalent to Eq. 2.9 may take advantage of the circular convolution theorem as follows

$$\mathrm{ST}_x[n, p_0] = \sum_{q=-N/2}^{N/2-1} X[q+p_0]W[q,p_0]\mathrm{e}^{j\frac{2\pi}{N}nq} = N \cdot \mathrm{DFT}_q^{-1}\Big\{X[q+p_0]W[q,p_0]\Big\} \tag{2.10}$$

where $X[q]$ and $W[q, p_0]$ are the DFTs of $x[n]$ and $w[n, p_0]$, respectively, and $\mathrm{DFT}^{-1}\{.\}$ is the inverse DFT (IDFT). The outcome of this alternative equation can be calculated for all frequencies by N FFT blocks of size N each. These blocks perform on the order of $O(N^2 log_2 N)$ operations, while using the same memory space as the implementation defined by Eq. 2.9. Thus, the more are the samples, the more efficient is the processing based on the FFT in comparison with the direct convolution. Nevertheless, the computational reduction comes at the cost of additional latency, as mentioned before for the STFT. This price might be excessive for real-time applications, especially if long time series are analyzed.

2.3 Undersampling the S-Transform

In general, TF representations count N^2 points out of an N-point signal. Such a redundancy is present in the ST as well, the complexity of which has a cubinc growth with the amount of points. Redundancy can be reduced by minimizing the densities of samples necessary to represent the signal with little or no aliasing in the TF plane. Following the considerations made in [33] regarding the STFT, we might likewise apply the Nyquist theorem on the set of windows in Eqs. 2.6 and 2.7 for the ST. By doing so, we may notice that the minimum sampling rate in the time domain for each voice varies with the voice itself: it is determined by the effective bandwidth, over which the frequency response of the respective window exceeds a certain threshold. Since the DFTs of the windows in Eq. 2.6 are also Gaussian, their poor selectivity in the frequency domain inevitably implies a certain amount of aliasing. The relevance of this effect depends on the threshold adopted to set the time sampling rate. Hereinafter, we approximate the Gaussian functions as their relative values within the frequency range defined by the full width at tenth maximum (FWTM) of the peak amplitude. This choice is motivated by the aim of limiting the amount of energy related to components aliased in frequency to the 10% of the total, at most. Therefore, given any non-zero frequency bin, the sampling rate for the respective voice, namely at the input of the Gaussian filter associated to it, may be bounded as

$$f_{s_p} \geq f_x \frac{\sqrt{2\ln(10)}}{\pi\sigma_p} = f_x \frac{|p|\sqrt{2\ln(10)}}{\pi N}, \qquad p \neq 0. \tag{2.11}$$

Here, the passband bandwidth is the FWTM of $|W[q, p]|$ that is twice the cut-off frequency at which the log-magnitude of $X[q]W[q, p]$ is at least 10 dB below the peak.

As for the zero-frequency voice, it holds that

$$f_{s_0} \geq \frac{f_x}{N}. \tag{2.12}$$

Since the tightest constraint on f_s is the following

$$f_x \, \frac{\sqrt{2\ln(10)}}{2\pi} \approx 0.34\, f_x \tag{2.13}$$

for $p = -N/2$, it is plain to see that setting $f_{s_p} = f_x$ means heavily oversampling all the voices. Nevertheless, for some voices, the input signal should be sampled at a rate higher than the Nyquist sampling rate in order to neglect the so-called *self-aliasing*. This undesirable effect occurs if the effective bandwidth of a voice encompasses folded portions of the periodic spectrum of $x[n]$. Assuming that f_x is actually coinciding with the Nyquist rate, we define the voices that suffer from self-aliasing adopting the same approximation used in Eq. 2.11 to be the ones that satisfy the following inequality

$$f_x \, \frac{|p|}{N} \left(1 + \frac{\sqrt{2\ln(10)}}{2\pi}\right) \approx 1.34 f_x \, \frac{|p|}{N} \geq \frac{f_x}{2} \tag{2.14}$$

which can be approximated as

$$\frac{|p|}{N} \geq 0.37. \tag{2.15}$$

Consequently, either the input rates for these voices are sufficiently increased above the Nyquist rate, or alternatively, these voices are discarded from the analysis by lowpass pre-filtering the input signal, so by removing the frequency components most self-aliased (i.e., with $|p| \geq 0.37N$). The factor in Eq. 2.14 is roughly equal to that mentioned in [23], which implicitly takes into account a 7.39-dB cut-off frequency. Self-aliasing affects only complex-valued signals, namely if the spectral energy is distributed over both the positive and negative sides of the TF representation, as shown for example in Fig. 2.3. In fact, it can be avoided for real-valued time series by resorting to the analytic form of the input signal, without the need for oversampling. As far as the minimum sampling rate in the frequency domain is concerned, it is supposed to be at least equal to the effective duration of the impulse response of the set of windows at every time instant. The sampling theorem cannot be strictly applied to the functions in Eq. 2.6, because they are frequency-dependent. Nevertheless, we might still infer a constraint by considering each of them separately, as if they were used in several STFTs. With such an expedient, the frequency-domain sampling rate can be bounded by

$$F_{s_p} \geq \frac{2\sqrt{2\ln(10)}\,\sigma_p}{f_x} = \frac{2N\sqrt{2\ln(10)}}{f_x\,|p|}, \qquad p \neq 0 \tag{2.16}$$

where the duration in samples is approximated as the FWHM of $w[n,p]$, according to the same criterion behind Eq. 2.11. In comparison to the initial sampling rate in frequency (i.e., N/f_x), we notice here that the voices associated to $|p| < 2\sqrt{2\ln(10)}$ do not satisfy Eq. 2.16, because their windows have significant values over time intervals longer than the duration of N samples. Yet, in practice, frequency sampling the ST at $p = -N/2, 1, ..., N/2-1$ is enough to avoid aliasing in time, because the width of any window is actually limited by the extension of $x[n]$. However, the Gaussian

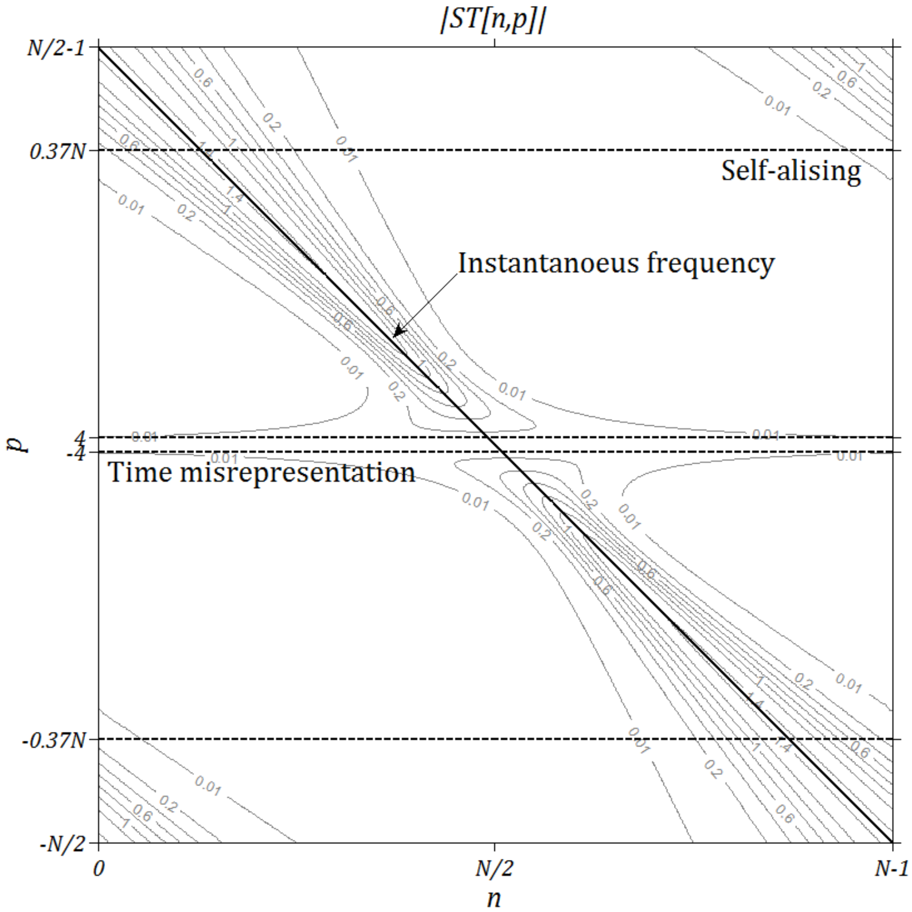

FIGURE 2.3: Amplitude of the ST for a complex chirp scanning the digital bandwidth (i.e., $[-N/2, N/2-1]$).

curves of the windows that do not comply with Eq. 2.16 are nearly flat within an interval of N samples, as we may appreciate in [36, Fig. 2]. This fact involves a detrimental effect: the components within a frequency range of $|p| < 2\sqrt{2\ln(10)} \approx 4.29$ are smoothed to an extent that their localization in time is severely undermined. The loss in the time resolution is actually complete at the zero frequency, where the rectangular window averages the samples over time. This *time misrepresentation* of the spectral components at the slow frequencies (i.e., around $p = 0$) creates ambiguity about the signal IF within a certain range around the frequency zero and is clear for example in Fig. 2.3. The symmetric boundaries for the relevance of the effects of self-aliasing and time misrepresentation according to Eqs. 2.15 and 2.16 are indicated in the figure, respectively. The bandwidth most affected can be shrunk in terms of hertz by either reducing f_s or increasing N, but cannot be eliminated. It is worth noting that this degradation does not occur in the frequency domain, because it would require the effective bandwidth of the window to exceed the maximum value imposed by N (i.e, $|p| > \pi N^2\sqrt{2\ln(10)} > N/2$).

Aliasing, self-aliasing, and time misrepresentation are side effects related to the construction of the ST according to the uncertainty principle. In the first place, undersampling always causes residual amount of aliasing, because the Gaussian windows are compactly-supported neither in time nor frequency. Secondly, self-aliasing and time misrepresentation exist because the temporal and spectral lengths of the windows are tied to the progressive TF resolution trade-off. It is important to understand that these intrinsic features of the ST are not flaws, because they do not prevent the transform from being invertible, as long as no modification is performed. However, their effects are harmful when the transform is to be used for synthesizing a modified TF representation or for estimating the signal IF, as we will see later on.

2.4 Signal Reconstruction

In many applications, TF analysis tools are used to provide an intermediate representation that is to be modified in order to change the characteristics of a nonstationarity signal. Of particular interest in this regard are TF excision techniques for jammer mitigation, such as those in [3] and [43] that are the subject of Chapter 4. After being applied, the modifications are made effective as soon as the a signal is reconstructed from the respective TF representation. In what follows, we describe how to efficiently retrieve the input time series from the output of the aforementioned digital filter bank. The synthesis has potentially a latency suitable to real-time applications. Furthermore, by resuming the initial comparison with the STFT, we may give interesting insights about the ST invertibility.

2.4.1 S-Transform Inversion

The original and discrete ST is always invertible. More specifically, in the absence of intermediate modifications, $\mathrm{ST}_x[n,p]$ is the unique TF representation of $x[n]$, which is then exactly recoverable from the respective transform. We emphasize here that invertibility follows from employing analysis windows that fulfill two generic and necessary conditions:

1. their sliding time responses span the whole time interval of N samples, during which the input signal is observed;

2. the composite frequency response of their modulated bandwidths cover the entire spectrum of the input signal with non-zero values.

These requirements basically state that the transform should preserve the information of the signal analyzed. The sufficient conditions instead are specific to the procedure used to invert the transform. As we anticipated, there are two main synthesis methods to map $\mathrm{ST}_x[n,p]$ back to $x[n]$ and they are both discussed in [36]. On the one hand, the FI exploits the linear relation between the ST and the FT: every frequency component of the Fourier spectrum results from summing over time the corresponding local components of the ST, according to

$$\frac{1}{N}\sum_{q=-N/2}^{N/2-1}\left(\sum_{m=0}^{N-1}\mathrm{ST}_x[m,q]\right)e^{j\frac{2\pi}{N}nq}=\frac{1}{N}\sum_{q=-N/2}^{N/2-1}X[q]e^{j\frac{2\pi}{N}nq}=\mathrm{DFT}_q^{-1}\{X[q]\}=x[n]. \tag{2.17}$$

The ST is so inverted by performing the inverse DFT of the one-dimensional function that is obtained by time-averaging the forward transform. On the other hand, the TI

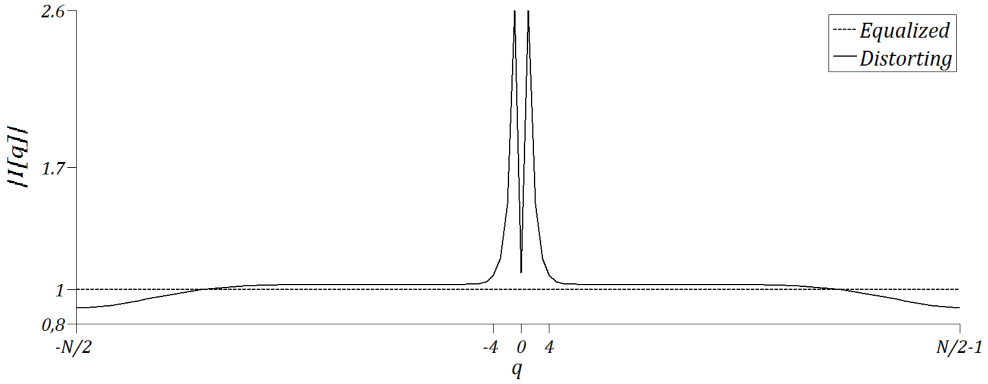

FIGURE 2.4: Amplitude of the composite transfer function of the bank of filters implementing the discrete ST.

approximates $x[n]$ as

$$\begin{aligned}\tilde{x}[n] &= \sum_{\substack{q=-N/2,\\ q\neq 0}}^{N/2-1} \frac{\sqrt{2\pi}}{|q|}\mathrm{ST}_x[n,q]\mathrm{e}^{j\frac{2\pi}{N}nq} + \mathrm{ST}_x[n,0] \\ &= \sum_{m=0}^{N-1} x[m]\left(\frac{1}{N}\sum_{q=-N/2}^{N/2-1} \mathrm{e}^{-\frac{1}{2}\left(\frac{(n-m)q}{N}\right)^2}\mathrm{e}^{j\frac{2\pi}{N}(n-m)q}\right) = x[n]\circledast i[n]\end{aligned} \tag{2.18}$$

where $i[n]$ denotes a smoothing function defined as the summation of all the windows normalized and frequency-shifted:

$$i[n] = \sum_{\substack{q=-N/2,\\ q\neq 0}}^{N/2-1} \frac{\sqrt{2\pi}}{|q|}w[n,q]\mathrm{e}^{j\frac{2\pi}{N}nq} + w[n,0] = \frac{1}{N}\sum_{q=-N/2}^{N/2-1} \mathrm{e}^{-\frac{1}{2}\left(\frac{nq}{N}\right)^2}\mathrm{e}^{j\frac{2\pi}{N}nq}. \tag{2.19}$$

the frequency response of which is denoted by $I[p]$ and depicted in Fig. 2.4. This function is real because it coincides the DFT of Eq. 2.19 that is symmetric. The exact input signal is therefore obtained by deconvolving $i[n]$ from the approximation $\hat{x}[n]$. Eqs. 2.17 and 2.18 interchange the roles of time and frequency under different sufficient conditions and computational requirements. We may assess these aspects of the ST by analogy with the STFT in [33]. Indeed, the complementarity between the TI and the FI reflects the duality of the filter-bank summation (FBS) and the overlap-add (OLA) methods used to invert the STFT, which both are reviewed in [38]. These two methods are motivated from the DFT and filter-bank interpretations of the STFT, respectively. In more detail, the OLA recovers $x[n]$ through IDFT the cross sections of $\mathrm{STFT}_x[n,p]$ taken for each n at which the analysis is performed, then summing the results over n. As such, the complexity of this procedure is dominated by that of the DFT, which is on the order of $O(N_w log_2 N_w)$ operations per sample if FFT blocks are used. Since the FI performs the same procedure but in the opposite order, it likewise performs on the order of $O(N log_2 N)$ operations per sample. The OLA ensures the STFT invertibility if the analysis window shifted by the samples at which the STFT is evaluated sums to the area under the window, or equivalently, if the sampling rate in time of the STFT is dense enough: at least twice the the cut-off frequency of the prototype lowpass filter used to perform the analysis. A generalized formulation of this constraint may be applied to frequency-dependent windows for the ST as follows

$$\sum_{n} w[n-m, p] = W[0, p]. \tag{2.20}$$

The summation above runs over the samples at which the ST analysis is performed. Satisfying the condition above is straightforward, and so the FI is exact, for the N^2-point ST. The FBS method synthesizes $x[n]$ by summing over frequency the products of $\mathrm{STFT}_x[n, p]$ and linear phase factors, which act as demodulators, with a complexity on the order of $O(N_w)$ per sample. The time series is exactly recovered from the STFT if the frequency responses of the passband filters used in the analysis add up to a constant. Equivalently, this implies choosing a window, which is either shorter than the frequency sampling factor (i.e., $N_w < N$) or larger and zero-valued at multiples of this factor, (i.e., $w[kN] = 0$ with $k \in \mathbb{Z}$). The constraint may be generalized for picking frequency-dependent windows according to

$$\frac{1}{N} \sum_{q} \frac{w[n, q]}{w[0, q]} e^{j \frac{2\pi}{N} nq} = \delta[n] \tag{2.21}$$

where $\delta[n]$ is the Kronecker delta function, and the summation runs over all the analyzed voices. Any fixed-resolution window for short-time analysis satisfies the above constraint if the number of frequency bins is at least equal to the window duration N_w, or equivalently, if the STFT is sampled in the frequency domain in accordance with the Nyquist theorem. This fact is exploited in [44] to build a vocoder, which is an analysis-synthesis system with limited data rate for speech signal processing. On the contrary, a set of windows for multi-resolution analysis can hardly comply with Eq. 2.21 without compromising the underlying TF resolution trade-off, unless it is specifically designed to do so. The Gaussian windows used in the original ST are the optimal choice for TF analysis, but they do not satisfy this constraint even when N voices are analyzed. Indeed, their modulated frequency responses make up a set of parallel bandpass filters, the composite transfer function of which preserves the energy but produces distortion. The distorted overall frequency response is plotted in [37, Fig. 2] and, after a proper energy normalization, in Fig. 2.4. The impulse response of this distortion is deduced from the first member of Eq. 2.21, which coincides with Eq. 2.19 when Eqs. 2.6 and 2.7 are adopted. It is now clear why, unlike the FBS method, the TI has to compensate for a smoothing function through the deconvolution in Eq. 2.18, at the cost of extra computations with respect to the nominal $O(N)$ operations per sample. In principle, the same issue would affect the FI, if the OLA method is applied to invert a version of Eq. 2.5 that employs time-dependent windows (e.g., with $v = m$) to achieve a time-varying TF resolution.

As opposed to [23], we define the discrete ST using the DFT of the Gaussian windows rather than discretizing the respective continuous FT. Moreover, we normalize the Gaussian windows given by Eqs. 2.6 and 2.7 to their actual discrete summations instead of $|p|/(\sqrt{2\pi}N)$, when constructing $i[n]$. These simple corrections are related to Eqs. 2.20 and 2.21. They preserve the energy and overcome the loss of information demonstrated in [36] at the slow frequencies. As a result, we can combine both the DFT and the filtering formulations of the ST with the TI or the FI, indiscriminately, without producing artifacts. We accomplish this achievement for instance in Fig. 2.5, by recovering the signal represented in Fig. 2.3. The plots of the samples returned by the TI and the FI coincide with the input time series to machine precision. As such, the ST exhibits the *neutrality* with respect of the forward transforms, if no intermediate modifications are performed. Interestingly, the side effects of filtering the

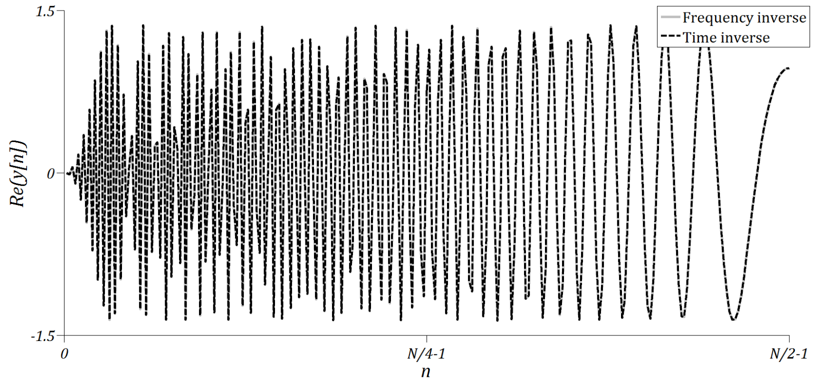

FIGURE 2.5: Real values of the time series denoted as $y[n]$ which exactly recovers the complex chirp $x[n]$ from the ST in Fig. 2.3.

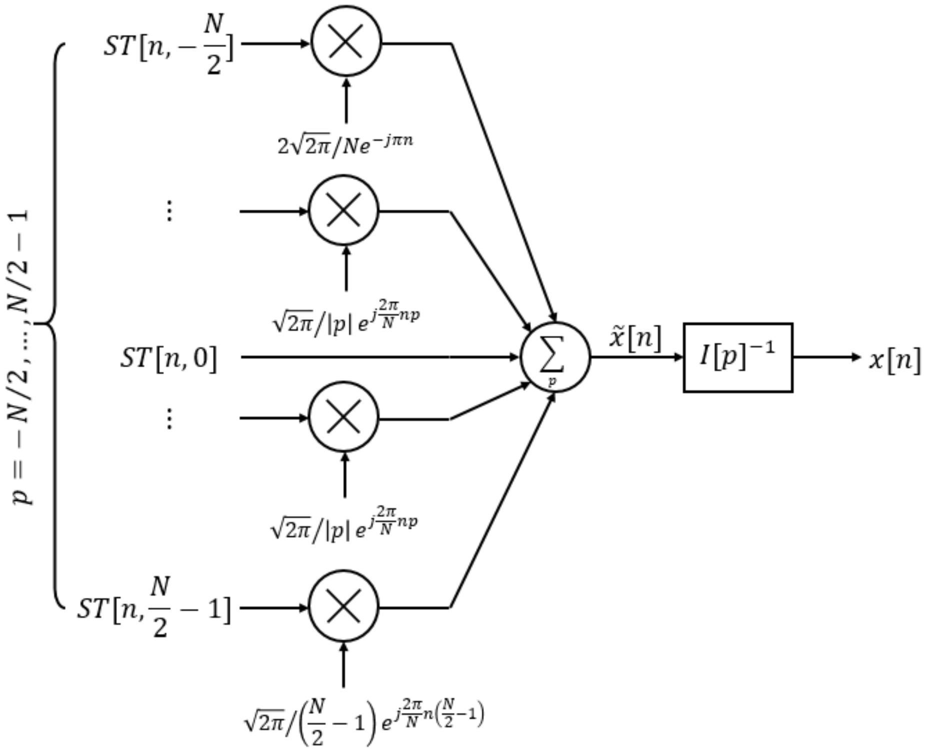

$ST[n, -\frac{N}{2}]$
$2\sqrt{2\pi}/Ne^{-j\pi n}$
$\sqrt{2\pi}/|p|\, e^{j\frac{2\pi}{N}np}$
$ST[n, 0]$
$\sqrt{2\pi}/|p|\, e^{j\frac{2\pi}{N}np}$
$ST[n, \frac{N}{2} - 1]$
$\sqrt{2\pi}/\left(\frac{N}{2} - 1\right) e^{j\frac{2\pi}{N}n\left(\frac{N}{2}-1\right)}$
$p = -N/2, \ldots, N/2 - 1$
$\sum_p$
$\tilde{x}[n]$
$I[p]^{-1}$
$x[n]$

FIGURE 2.6: Reconstruction of a time series by means of the TI.

ST described in [36] resemble those for the STFT in [33]. On the one hand, the synthesis of the TI is smoothed by the convolution between the input and the IDFT of the window-weighted version of any spectral modification. On the other hand, the output synthesized through the FI is smeared due to an extra undesired windowing operation. Consequently, the choice between the TI and the FI maintains either the time or the frequency localization of the implemented TF filter, respectively. If the aim is reconstructing one sample per unit time, as a data stream, the TI is the natural for recovering the input signal. Therefore, the samples at the output of each channel in Eq. 2.9 are first multiplied by a frequency-dependent factor, then accumulated, and finally deconvolved through Eq. 2.19. This latter operation may be interpreted as that of an equalization filter, which reverses the distortion mentioned before. In the light of the formulation of the analysis and inversion with the ST, the whole real-time implementation consists of a bank of digital filters followed by the block diagram in Fig. 2.6. In the absence of modifications, it is a linear and time-invariant all-pass system with unitary gain. Furthermore, the TI offers a couple of advantages over the FI. First of all, synthesis is possible even if the ST is processed through a subset of channels. This possibility is useful to restrict the analysis to the voices of interest, as we do in [3]. While it is not possible with the FI, since it requires evaluating the ST for as many frequency voices as the signal data samples. This requirement is a very demanding obligation for some applications, as noted in [36]. Secondly, any modification changes instantaneously the signal synthesized, because the time localization of the TF filter is maintained. This fact prevents aliasing in time, which instead could emerge with the FI, due to the enlarged duration of the filter impulse response.

2.4.2 Instantaneous-Frequency Estimation

Instead of inverting the ST, a mono-component signal can be reconstructed, with limited reconstruction error, by estimating the respective IF. The publication in [45] reviews the interpretation and the estimation of the IF for discrete and non-stationary signal. On this regard, one simple and widely used approach for IF estimation extrapolates the energy peak from a neat TF representation. This method capitalizes on the fact that the energy distribution in the TF domain should be concentrated around IF trajectory. Accordingly, we assign the following frequency as an IF estimate:

$$\hat{f}[n] = \arg\max_{p} |\mathrm{ST}_x\,[n,p]|^2. \tag{2.22}$$

The same approach but using other TF representations was adopted in [46–48] to evaluate the respective estimation performance, which is beyond the scope of the chapter. As emphasized before, the ST has a direct relation with the FT. Namely, if the following sinusoid

$$x[n] = A\,\mathrm{e}^{j(2\pi f_0 n + \phi_0)} \tag{2.23}$$

is analyzed through the ST, as argued in [32], we have that

$$\mathrm{ST}_x[n, f_0] = A\,\mathrm{e}^{j\phi_0} \tag{2.24}$$

where A is the amplitude, f_0 is the sinusoid frequency, and ϕ_0 is a constant phase shift. Let us consider a non-stationary signal with constant or slowly varying amplitude

$$x[n] = A\,\mathrm{e}^{j\phi[n]} \tag{2.25}$$

where $\phi[n]$ represents the discrete-time instantaneous phase. The corresponding discrete-time IF is given by

$$f[n] = \frac{1}{2\pi}\frac{\partial}{\partial n}\phi[n]. \tag{2.26}$$

Therefore, $\phi[n]$ can be written according to

$$\phi[n] = 2\pi f[n]\,n + \theta[n] \tag{2.27}$$

where $\theta[n]$ is a time-dependent factor. Now, by generalizing Eq. 2.24, we see that the ST at a specific time instant and at the corresponding IF returns

$$\mathrm{ST}_x[n, f[n]] = A\,\mathrm{e}^{j\theta[n]}. \tag{2.28}$$

Consequently, a non-stationary signal can be reconstructed by first estimating the IF following Eq. 2.22, then extracting approximations to the amplitude and instantaneous phase using Eqs. 2.27 and 2.28. The signal is finally recovered as

$$\hat{x}[n] = \hat{A}[n]\, \mathrm{e}^{j(2\pi \hat{f}[n]\, n + \hat{\theta}[n])} \tag{2.29}$$

where

$$\mathrm{ST}_x[n, \hat{f}[n]] = \hat{A}[n]\, \mathrm{e}^{j\hat{\theta}[n]}. \tag{2.30}$$

The one described above is a novel procedure to a reconstruct non-stationary signal from the ST. Now, we can employ this method for inferring a time series both directly and indirectly. Under our assumption of constant-amplitude waveform, $x[n]$ can be directly estimated by means of Eqs. 2.29 and 2.30. However, to account for IF ambiguity around the frequency zero and to enable the signal reconstruction when fewer voices than the total N are analyzed, we choose to average the amplitude estimates over the observation time, according to

$$\hat{A} = \frac{1}{N} \sum_{m=0}^{N-1} \hat{A}[m]. \tag{2.31}$$

Alternatively, the IF can be estimated as the polynomial curve that interpolates the peaks of $|ST_x\,[n,p]|^2$. This second procedure, however, requires an a-priori model for the IF trajectory. This model is compatible neither with the effect of the IF ambiguity nor with the imposition of a restricted number of voices. Thence, it eventually introduces significant errors in the estimation under common circumstances. Both the two reconstructions assume a single frequency-modulated waveform, hence a mono-component signal, so that there is only one energy peak per time unit. This assumption is a actually simplification. It can be relaxed if this procedure is expanded to address a multi-component signal made by the superposition of multiple waveforms, which might be either well separated or crossing each other in TF plane. Such an upgrade is not of interest, because it is generally addressed by adding tracking capabilities to the IF estimator, which is beyond the scope of the work described. In any case, the waveforms are discriminated from the white background noise by revealing the TF components that exceed a frequency-dependent energy threshold. In the following as well as Chapters 3 and 4, the detection is performed at constant false-alarm rate by leveraging of the additive and Gaussian nature of the noise.

2.5 Complexity-Scalable S-Transform

Containing the ST complexity is the key to expand its utility to a broader field of application. For instance, the computational burden is particularly challenging when this transform is implemented as a bank of filters able to work in near real time, hence introducing a latency that is only related to the filers themselves. The complexity-scalable sampling scheme illustrated below computes the forward ST with controllable efficiency, generalizing the schemes proposed in [20] and [32]. This flexibility is achievable as a combination between *temporal decimation* and *spectral compression*.

2.5.1 Decimating the S-Transform in Time

The original ST produces a TF representation of N^2 points out of an N-point signal. The size is the same as for the fully-redundant STFT given by Eq. 2.1, which can be effectively decimated in the time and frequency domains, as described in [33]. Likewise, the ST can be undersampled, according to the sampling theorem, only in the time domain, due to the intrinsic dependency of the windows on frequency. To generalize the criterion underlying Eq. 2.11, we introduce a dimensionless and positive parameter denoted by α to tune the effective width of the DFT of the window in Eqs. 2.6 and 2.7 that is approximated as the relative values within a bandwidth limited by a cut-off frequency of arbitrary attenuation of $10\log_{10}(\alpha)$ in dB, as shown in Fig. 2.8. First, we define the following discrete and piecewise function with respect to α:

$$L_k(\alpha) = \begin{cases} N, & |k| \geq \lceil \frac{\pi N}{\sqrt{2\ln\alpha}} \rceil \\ 2^{\lfloor \frac{|k|\sqrt{2\ln(\alpha)}}{2\pi} \rfloor}, & \lceil \frac{2\pi}{\sqrt{2\ln\alpha}} \rceil \leq |k| < \lfloor \frac{\pi N}{\sqrt{2\ln\alpha}} \rfloor \\ 1, & |k| < \lfloor \frac{2\pi}{\sqrt{2\ln\alpha}} \rfloor \end{cases} \tag{2.32}$$

with $k \in \mathbb{Z}$ and $\alpha > 1$. Then, we can provide a generalized formulation of the effective bandwidth of each voice in terms of samples by restricting the countable and infinite domain of $L_k(\alpha)$ to the finite range of frequencies under analysis (i.e. $k = p \in \{-N/2, ..., N/2-1\}$). The formulation in Eq. 2.32 describes curves depicted in Fig. 2.7, and implies that

$$L_p(\alpha) = 1, \forall p \Leftrightarrow 1 < \alpha < \exp\left(\frac{8\pi^2}{N^2}\right) \tag{2.33}$$

which corresponds to the maximum temporal decimation that converts the ST into the FT. On the contrary, the decimation is practically ineffective if

$$L_p(\alpha) = N, \forall p \neq 0 \Leftrightarrow \alpha \geq \exp\left(\frac{\pi^2 N^2}{2}\right). \tag{2.34}$$

Between these two intervals, the greater is α, the larger is $L_p(\alpha)$, thus the better is the approximation of the windows, particularly for high-rate voices.

$$\widetilde{W}[q,p] \approx W[q_p,p], \quad |p| \geq \lceil \frac{2\pi}{\sqrt{2\ln\alpha}} \rceil \tag{2.35}$$

for $q_p = -L_p(\alpha)/2, 1, ..., L_p(\alpha)/2 - 1$. Since the windows at lower $|p|$ are nearly flat, their respective DFTs are well approximated by a single sample as

$$\widetilde{W}[q,p] \approx W[0,p], \quad |p| < \lfloor \frac{2\pi}{\sqrt{2\ln\alpha}} \rfloor \tag{2.36}$$

with $L_p = 1$. This latter approximation is valid for the the window in Eq. 2.7, regardless of α, which means that the zero-frequency voice can be always undersampled. Along the same lines, we can define the following frequency-dependent temporal decimation factor

$$D_p(\alpha) = \frac{N}{L_p(\alpha)} \tag{2.37}$$

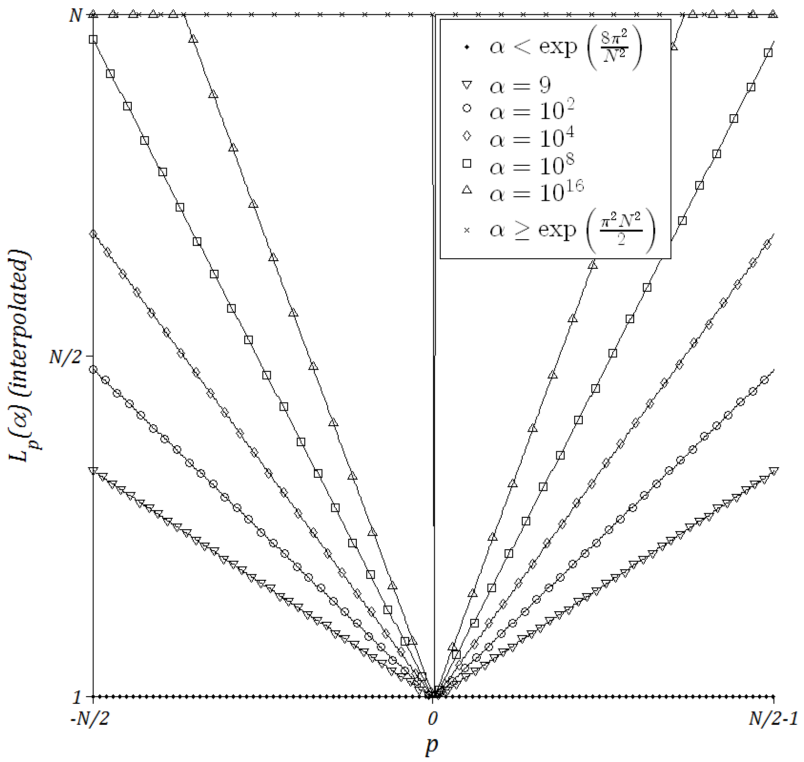

FIGURE 2.7: Progression of the effective bandwidth of the ST frequency-dependent window (the markers are there just to distinguish the curves for different values of α).

that is determined by α for the windows associated to $|p| \geq \lceil 2\pi/\sqrt{2\ln\alpha}\rceil$, and reduces to N for those associated to $|p| < \lfloor 2\pi/\sqrt{2\ln\alpha}\rfloor$. Consequently, the Nyquist sampling theorem can be applied now to formulate the time-decimated ST defined by

$$\mathrm{ST}_x[l_p(\alpha), p] = \sum_{m=0}^{L_p(\alpha)-1} x[m]\,\widetilde{w}[l_p(\alpha) - m, p]\,\mathrm{e}^{-j\frac{2\pi}{N}pm} \tag{2.38}$$

where $l_p(\alpha) = \lfloor n_p\, D_p(\alpha)\rfloor$, $n_p = 0, 1, ..., L_p(\alpha) - 1$, and $\widetilde{w}[m,p]$ denotes the truncated window given by the IDFT of the following frequency response

$$\widetilde{W}[q,p] = \begin{cases} W[q,p], & \lfloor\frac{-L_p(\alpha)}{2}\rfloor \leq q \leq \lceil\frac{L_p(\alpha)}{2}\rceil - 1 \\ 0, & \text{otherwise.} \end{cases} \tag{2.39}$$

Accordingly, we see that the sampling rate in the time domain depends on the analyzed voice, as it changes with the frequency index p according to $f_x/D_p(\alpha)$. For a generic voice with index p_0, it holds that

$$ST_x[l_{p_0}, p_0] = L_{p_0}(\alpha) \cdot \mathrm{DFT}_q^{-1}\{X[q]\widetilde{W}[q,p_0]\}. \tag{2.40}$$

Any decimation factor corresponding to $\alpha \geq 10$ complies with the bound on the time sampling rate set in Eq. 2.11. Instead, choosing $\alpha \ll 10$ boosts the temporal decimation, but it changes severely the widows so that they are Gaussian neither in the frequency domain nor in time domain, and no longer of minimum time-bandwidth product. This choice is particularly detrimental for fast-rate voices, considering that

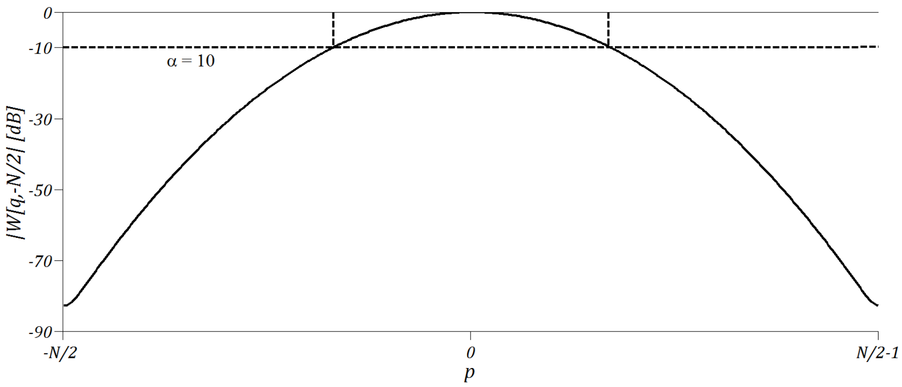

FIGURE 2.8: Amplitude of the transfer function of the least selective Gaussian filter (i.e., for $p = -N/2$).

when α is much smaller than 10, the spectral responses of their localizing windows confined within the effective bandwidth are nearly flat, which results in ringing effect in the time domain. However, regardless of the value selected for α, the input signal is exactly recovered through the FI of the transform in Eq. 2.38. Indeed, the constraint given by Eq. 2.20 reformulated with respect to $D_p(\alpha)$ is the following

$$\sum_{n_p=0}^{L_p(\alpha)-1} \widetilde{w}[\lfloor n_p\, D_p(\alpha)\rfloor - m, p] = \widetilde{W}[0, p] \qquad \forall\alpha. \tag{2.41}$$

The identity above is verified for any value of α besides a discretization error, which arises for the voices down-sampled by a non-integer $D_p(\alpha)$. As far as the TI is concerned, sampling the ST in the time domain according to Eq. 2.38, extends the synthesis procedure to include interpolation by $(D_p(\alpha))^{-1}$ for every undersampled voice in order to restore the decimation factor to unity before the inversion. This addition comes at the cost of an increase of complexity in terms of operations, which compensates what is saved in terms of memory elements. Moreover, the interpolation inevitably introduces errors, which reduce with increasing α. On the contrary, the efficiency of the FI benefits from the time decimation, because it can process the result of Eq. 2.38 directly.

Once applied to the ST filter-bank architecture, temporal decimation implies down-sampling by $D_{p_0}(\alpha)$ the time series at the inputs of the channels in Fig. 2.2. Therefore, the transform in 2.38 is implemented by a bank of N filters with $N/2 + 1$ different input rates, at most. In order to enable the TI, N interpolators are added at the input of the filter-bank summation and equalization depicted in Fig. 2.6. Real-time TF analysis and synthesis based on the time-decimated ST is then carried out through a *multi-rate* design.

2.5.2 Compressing the S-Transform in Frequency

We have argued that original ST is uniformly sampled in frequency and cannot be strictly decimated in this dimension. Nevertheless, the constant spacing between the voices entails unnecessary amount of spectral redundancy, because it is not consistent with the progressive scaling of the ST frequency-dependent windows. To compress the TF plane over the frequency domain, we may device a non-uniform sampling rate, which mirrors the uncertainty principle. The idea is to sample finely the slow frequencies, which are better localized in this domain, and more coarsely the fast frequencies, where the spectral resolution is poorer. Going into detail, we firstly define the following continuous function:

$$\gamma(\beta) = 1 + \frac{1}{\beta} \tag{2.42}$$

with respect to a dimensionless and positive parameter denoted by β such that $1 \leq \beta \leq N/2$, which results in $1 + 2/N \leq \gamma(\beta) \leq 2$. Let us then consider, for the sake of argument, the positive semi-axis of a generic and continuous variable, i.e., $\nu \in]0, +\infty[$. Now, we may use Eq. 2.42 to indefinitely segment this domain into contiguous partitions of increasing sizes along ν, which are delimited by $(\gamma(\beta))^k$. The increment at every step k is determined by β and at most doubles the size of the previous partition. The midpoint of every partition may be described through the following discrete-defined exponential function

$$\epsilon_k(\beta) = \frac{\gamma(\beta)^{k-1}(1 + \gamma(\beta))}{2} > 0 \qquad \forall k \in \mathbb{Z}. \tag{2.43}$$

We restrict this countable and infinite domain into a finite set such that $k \in \{k_1(\beta), ..., k_2(\beta)\}$, by assuming a pair of conditions. In the first place, the distance between consecutive midpoints should be at least unitary:

$$\epsilon_{k+1}(\beta) - \epsilon_k(\beta) \geq 1 \Leftrightarrow k_1(\beta) = \left\lceil \frac{\log(\frac{2}{\gamma(\beta)^2 - 1})}{\log(\gamma(\beta))} \right\rceil + 1 \tag{2.44}$$

where it is straightforward to prove that

$$\epsilon_{k_1}(\beta) > \beta. \tag{2.45}$$

Secondly, we bind the maximum value to

$$\epsilon_k(\beta) < \frac{N}{2} \Leftrightarrow k_2(\beta) = \left\lfloor \frac{\log(\frac{N}{1+\gamma(\beta)})}{\log(\gamma(\beta))} \right\rfloor + 1. \tag{2.46}$$

Finally, we can formulate a discrete and piecewise function based on Eqs. 2.43-2.46 in order to identify M frequency indices out of the N total, as follows

$$p_k(\beta) = \begin{cases} -\lfloor \epsilon_{|k|+1-\beta+k_1(\beta)}(\beta) \rfloor - 1, & k = -M/2, ..., -\lfloor \beta \rfloor - 1 \\ k, & k = -\lfloor \beta \rfloor, ..., \lfloor \beta \rfloor - 1 \\ \lfloor \epsilon_{k-\beta+k_1(\beta)}(\beta) \rfloor, & k = \lfloor \beta \rfloor, ..., M/2 - 1 \end{cases} \tag{2.47}$$

where $M = 2(\lfloor \beta \rfloor - k_1(\beta) + k_2(\beta) + 1)$ is always even. By sampling the frequency domain through Eq. 2.47, we obtain a non-uniform density of spectral components spread across the bandwidth according to a two-step behavior, which reflects the uncertainty principle: the sampling step is constant at low frequencies and grows nonlinearly as the frequency increases above a certain value. The progression of the sampling step in this domain is controlled by β, as shown in Fig. 2.9. Even though we refer to sampling rates, the function in Eq. 2.47 does not rigorously describe a spectral decimation. Therefore, we find more appropriate referring to this sampling scheme as compression, which reduces the order of operations and memory elements required for both analysis and synthesis by the ratio $N/M(\beta)$. Frequency

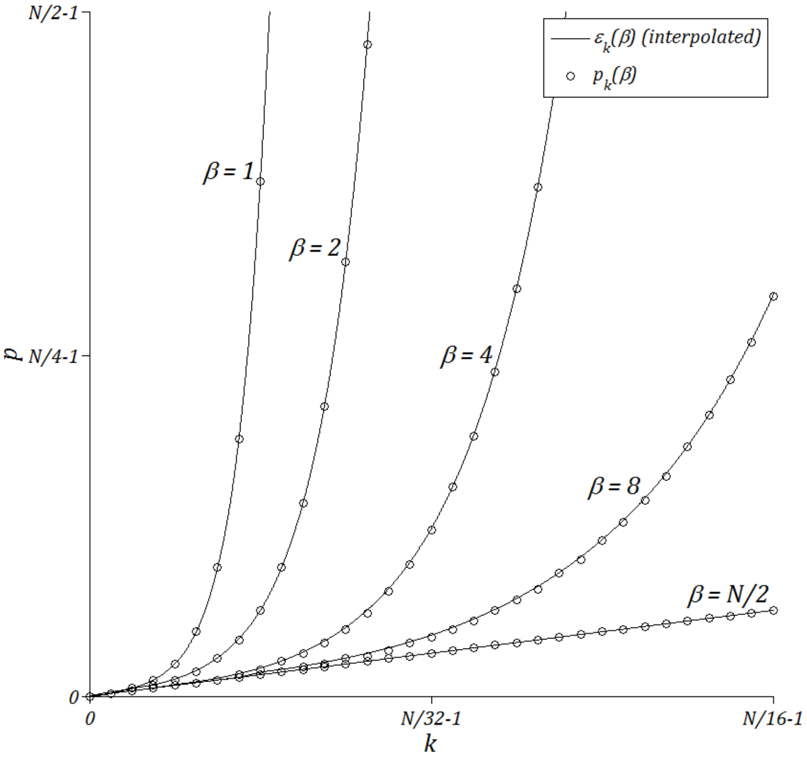

FIGURE 2.9: Progression of the frequency-domain sampling step after spectral compression, where the curve for $\beta = N/2$ corresponds to the uniform and minimum frequency spacing (the makers pinpoint the indices of the frequency bins sampled and denoted by $p_k(\beta)$).

compression enables downsizing the bank of filters computing the ST from the initial N channels to $M(\beta)$. The higher the compression ratio is the more distorted is the composite transfer function of this implementation. While the FI by definition cannot be employed with less than N voices, the additional distortion due to the reduction in the number of voices does not hinder the adoption or the performance of the TI, which in fact employs an equalization. In other words, as done for the N^2-point ST, the times series can be exactly recovered by equalizing the output of the filter-bank summation in Fig. 2.6 through the DFT of

$$i[n] = \frac{1}{N} \sum_{q=-M(\beta)/2}^{M(\beta)/2-1} \frac{w[n,q]}{w[0,q]} e^{j\frac{2\pi}{N}nq}. \tag{2.48}$$

The sole constraint on the choice of β is specified by the necessary conditions for invertibility mentioned in the previous section. Namely, the composite frequency response of the filter bank depicted in Fig. 2.2 for the voices $p_k(\beta)$ must cover the spectrum under analysis with nonzero values. Furthermore, since $p_k(\beta)$ is the central frequency of the pth Gaussian window given by Eq. 2.6, it also determines the window effective bandwidth, as defined using Eq. 2.32, when the ST is undersampled in time according to a specific α. Consequently, fulfilling the aforementioned condition constrains both the temporal decimation and the spectral compression according to

$$\begin{cases} \frac{L_{p_{k-1}}(\alpha)}{2} + \frac{L_{p_k}(\alpha)}{2} \geq |p_k(\beta) - p_{k-1}(\beta)| \\ \frac{L_{p_{M/2-1}}(\alpha)}{2} \geq \frac{N}{2} - 1 - p_{M/2-1}(\beta) \\ \frac{L_{p_{-M/2}}(\alpha)}{2} \geq \frac{N}{2} - |p_{-M/2}(\beta)| \end{cases} \tag{2.49}$$

with $k = -M/2, ..., -\lfloor\beta\rfloor - 1, \lfloor\beta\rfloor, ..., M/2 - 1$. Elaborating on the first inequality of Eq. 2.49, the following relationship between α and β can be obtained:

$$\alpha > \exp\left(\left(\frac{\sqrt{2}\pi}{1 + 2\beta}\right)^2\right). \tag{2.50}$$

Without spectral compression, i.e., for $\beta = N/2$, the above inequality implies that $\alpha > \exp\left(2\pi^2/(1+N)^2\right)$, consistently with (2.33). Instead, the maximum compression, i.e. for $\beta = 1$, requires, according to Eq. 2.50, that $\alpha > 8.96$. If this constraint is not fulfilled, information is not preserved and the TI cannot work.

2.5.3 Time-frequency Sampling Scheme

In the following, we refer to the combination of decimation in time and compression in frequency as a generic TF sampling scheme. Any scheme builds a ST matrix composed of the following number of elements

$$L(\alpha, \beta) = \sum_{p_k = -M(\beta)/2}^{M(\beta)/2-1} L_{p_k}(\alpha) \leq N(N-1) + 1 \qquad \forall \alpha, \beta. \tag{2.51}$$

This number is quantified by α and β, to which are proportional the necessary memory space and computational power. Note that because the zero-frequency voice can be always undersampled, $L(\alpha, \beta)$ is always less than N^2. As a result, the complexity of the time-decimated and frequency-compressed ST is ultimately scalable through α and β. The numbers of points obtained with different combinations of these parameters are depicted for N = 512 in Fig. 2.10. The curves are normalized, so that their trend is valid regardless of N. Besides illustrating the computational saving in the forward ST, this figure also highlights another important implication of the suggested TF sampling scheme, namely which synthesis method can be paired with it. Without compression, the FI can be performed with any decimation factor, up to averaging the representation over time into the one-dimensional N-point Fourier spectrum, i.e., $L(\alpha, \beta) = N$ with $\beta = N/2$. Likewise, the TI works perfectly with any spectral compression ratio; it can be used also with any time decimation, but only by interpolating the voices undersampled, as earlier mentioned. As far as the no inversion area plotted in Fig. 2.10 is concerned, it roughly indicates the combinations of α and β that do not satisfy Eq. 2.50. The foundations of the proposed TF sampling scheme are laid by the principles underlying the FST and the earlier DOST. Both these representations make use of an orthonormal basis functions constructed with a reduced set of frequency-dependent windows, which are local in time, and have compact and contiguous bandpass bandwidths; the resultant one-to-one representation coincides with that obtained by minimally sampling the TF plane with $\alpha = 9$ and $\beta = 1$, which are derived from Eq. 2.50. According to this combination, the voices are undersampled and selected such that they have adjacent and non-overlapping sub-bands covering completely the signal spectrum. Therefore, the FST

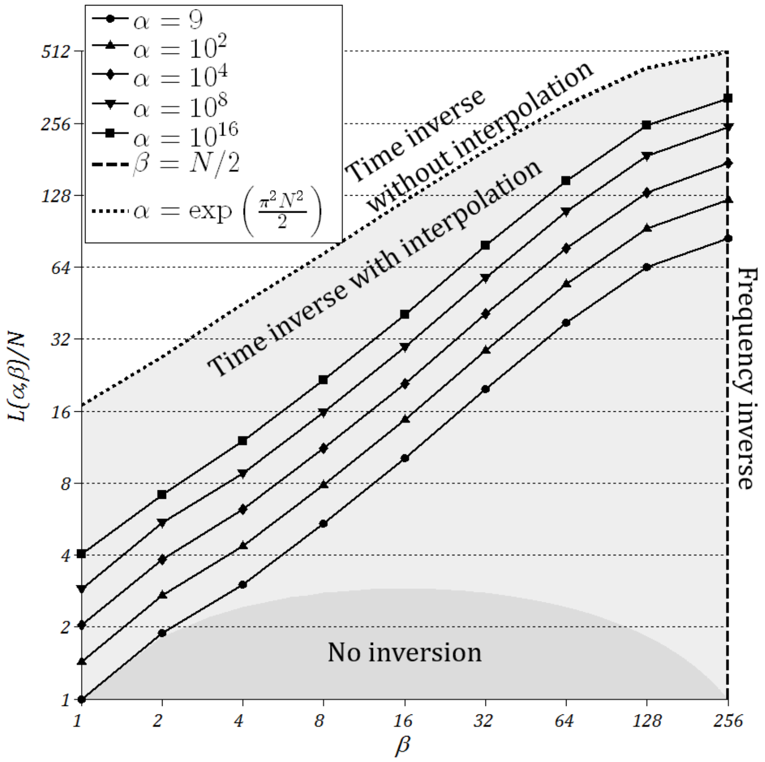

FIGURE 2.10: Progression of $L(\alpha, \beta)$ normalized on N (the dashed and the dotted curves correspond to no and maximum decimation in time, respectively).

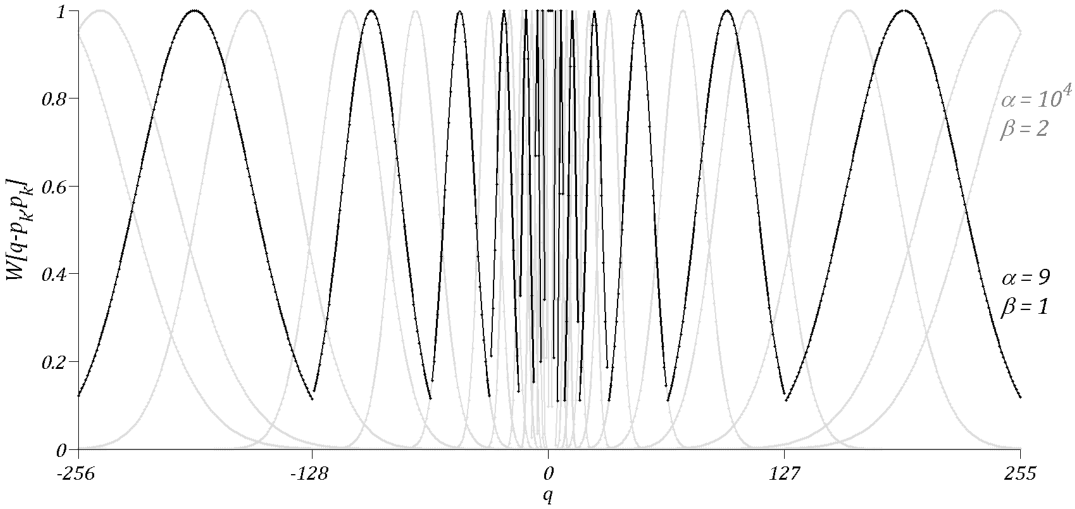

FIGURE 2.11: Superposition of two sets of windows for TF sampling schemes entailing different amounts of redundancy.

and the DOTS are covered by the more general and flexible framework of the proposed complexity-scalable ST, which offers different trade-offs between efficiency and accuracy. For the sake of clarity, the set of windows of the octave sampling scheme is shown together with that for $\alpha = 10^4$ and $\beta = 2$ in Fig. 2.11.
Eventually, it is worth briefly reviewing the algorithm for the FST presented in [20]. This forward transform is computed by shifting the DFT of the localizing window, rather than the DFT of the signal, such as in Eq. 2.10. As a consequence, this formulation is not mathematically equivalent to the original definition of the ST in Eq.

2.8, because it does not account for a phase factor, as becomes clear by elaborating on Eq. 2.9:

$$\begin{aligned} ST_x[n,p_0] &= \left(x[n]\mathrm{e}^{-j\frac{2\pi}{N}p_0 n}\right) \circledast w[n,p_0] \\ &= \left(x[n] \circledast w[n,p_0]\mathrm{e}^{j\frac{2\pi}{N}p_0 n}\right)\mathrm{e}^{-j\frac{2\pi}{N}p_0 n} \\ &= N\cdot \mathrm{DFT}_q^{-1}\{X[q]W[q,p_0]\}\cdot \mathrm{e}^{-j\frac{2\pi}{N}p_0 n} \end{aligned} \tag{2.52}$$

where the inverse DFT of the product between $X[q]$ and $W[q-p_0,p_0]$ undergoes a phase shift. Likewise, the same factor is missing in the inverse FST, which simply reverses the steps of the forward algorithm. Therefore, we can invert the octave scheme underlying the FST by reversing the steps in Eq. 2.52, thus also including the phase correction. This particular reconstruction procedure is referred to as *octave reverse* hereinafter.

2.6 Case Study: Time-Frequency Filter

In the present section, we examine a case study demonstrating TF filtering through the complexity-scalable ST. Our ultimate aim is evaluating the improvement in terms of reconstruction accuracies that is achievable by allowing for increasing degrees of redundancy in the TF representation. The synthesis is performed through the two reconstruction procedures based on either inverting the ST or IF estimation, in addition to the octave reverse specific for the minimal sampling scheme. In the scenario under study, the following frequency-modulated waveform is considered

$$s[n] = \exp\left(j(2\pi(f_0 + g\,n)\,n + \phi_0)\right) \tag{2.53}$$

the frequency rate of which is given by

$$g = \frac{0.37 - 20/N}{2(N-1)} \tag{2.54}$$

and it is shifted in phase by $\phi_0 = \pi/4$, with initial frequency $f_0 = 20/N$. Eq. 2.53 describes a complex-valued linear chirp sampled with a unitary sampling step, and is designed to scan the positive frequencies in the range $p \in [20/N, 0.37]$, hence conveniently avoiding the bandwidths heavily affected by IF ambiguity and self-aliasing, both explained before. The signal analyzed is expressed as

$$x[n] = s[n] + d[n] \tag{2.55}$$

with $N = 512$; it is the superposition of the waveform given by Eq. 2.53 and the following pair of bursts

$$d[n] = \begin{cases} \mathrm{e}^{j2\pi\frac{96n}{N}}, & n = 64, ..., 159 \\ \mathrm{e}^{j2\pi\frac{64n}{N}}, & n = 224, ..., 319 \\ 0, & \text{otherwise.} \end{cases} \tag{2.56}$$

The task is erasing the bursts smearing the chirp, before retrieving the filtered signal. For this purpose, a TF filter is put in place in the form of binary mask, which is multiplied point by point with the matrix computed through the forward ST. For the sake of simplicity, the employed mask blanks ideally the TF components containing more

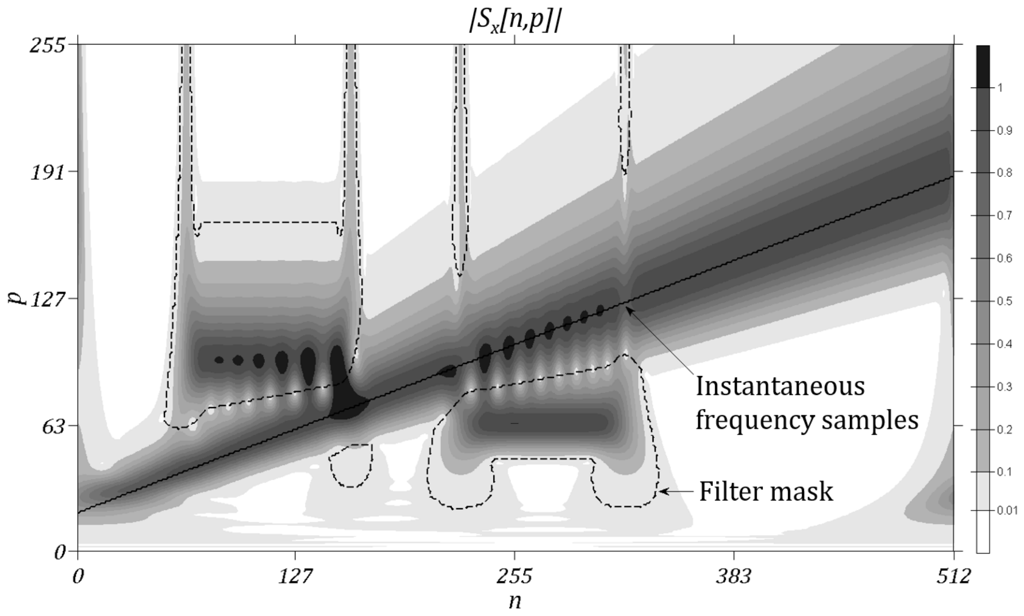

FIGURE 2.12: Amplitude the original ST of the signal given by Eq. 2.55. This TF representation coincides with that of the maximal sampling scheme for $\alpha = \exp\left(\pi^2 N^2/2\right)$ and $\beta = N/2$, except for the zero-frequency voice, which is undersampled for any combination of α and β.

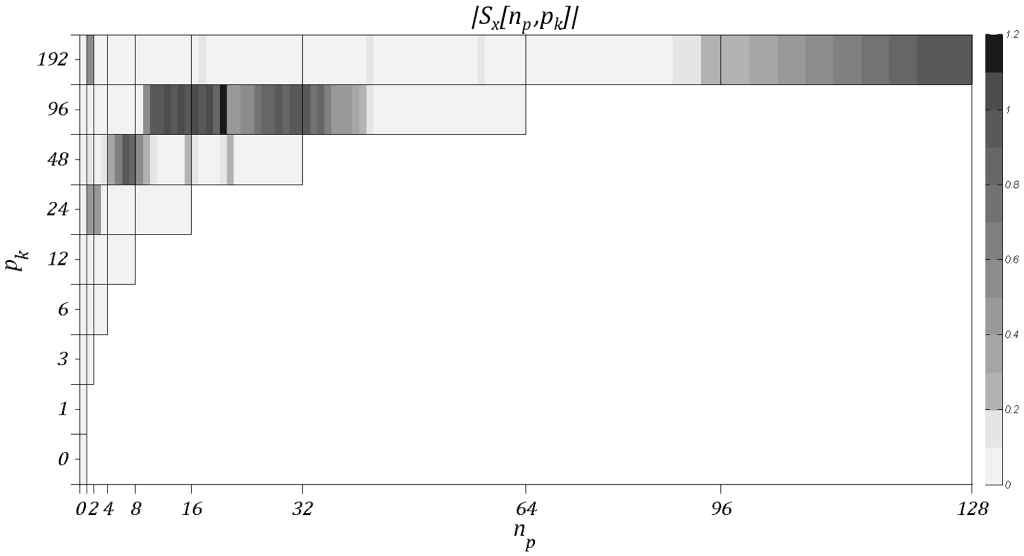

FIGURE 2.13: Amplitude the FST of Eq. 2.55, which coincides with the TF representation sampled for $\alpha = 9$ and $\beta = 1$.

energy from $d[n]$ than from $s[n]$. This simplification is meant to provide a common framework for the energy detection and is not a strong one, since is the comparison between the performance evaluated for STs of different redundancy that matters. The amplitude of the ST together with the respective filter mask is shown in Figs. 2.12-2.14 for the fully redundant ST, and the one-to-one FST. The performance for this case study is evaluated in terms of complexity reduction versus reconstruction accuracy. The accuracy is expressed in terms of the normalized root-mean-square

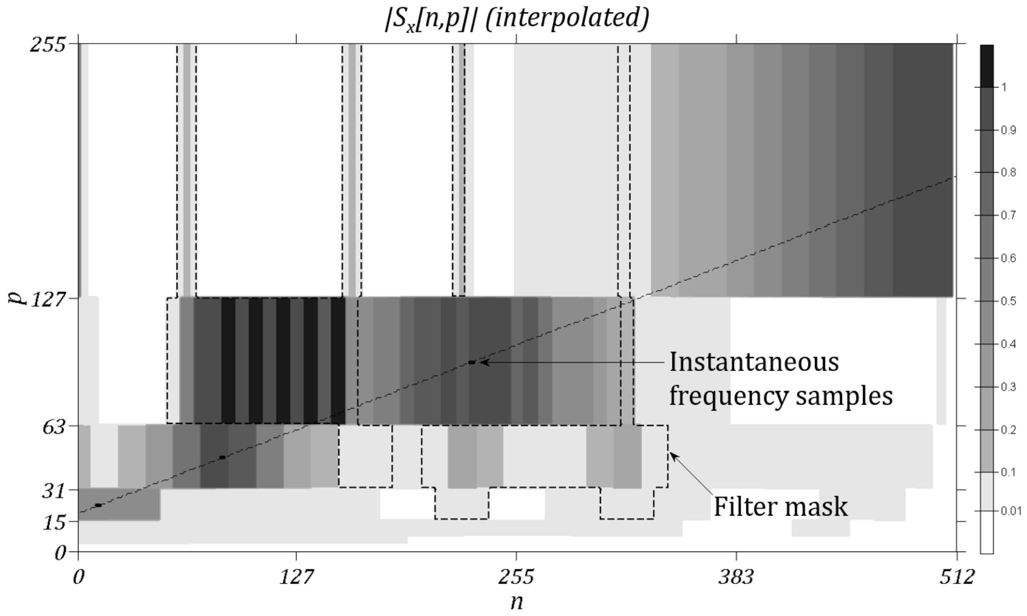

FIGURE 2.14: Amplitude of the nearest-neighbor interpolation of the TF representation in Fig. 2.13 as a comparison with that in Fig. 2.12.

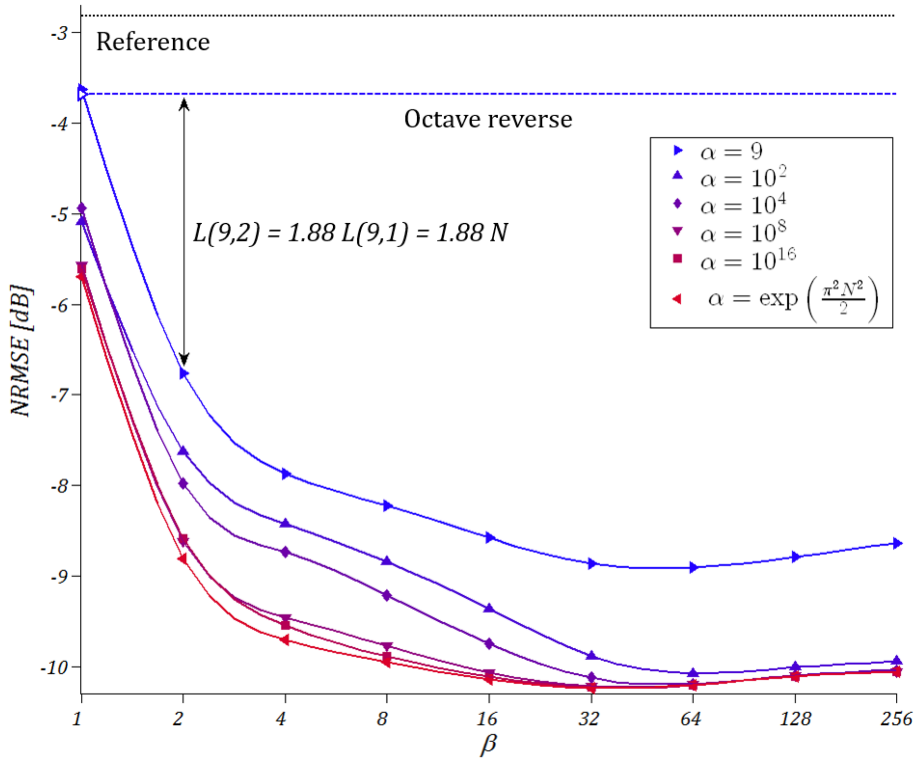

FIGURE 2.15: Reconstruction accuracy in NRMSE of the filtered signal synthesized through the TI for different pairs of α and β, in comparison with that obtained by means of the octave reverse, and with the reference error in dotted line.

error (NRMSE) defined by

$$\mathrm{NRMSE} = \left(\frac{\sum_{n=0}^{N-1} |\varepsilon[n]|^2}{\sum_{n=0}^{N-1} |y[n]|^2}\right)^{\frac{1}{2}} \tag{2.57}$$

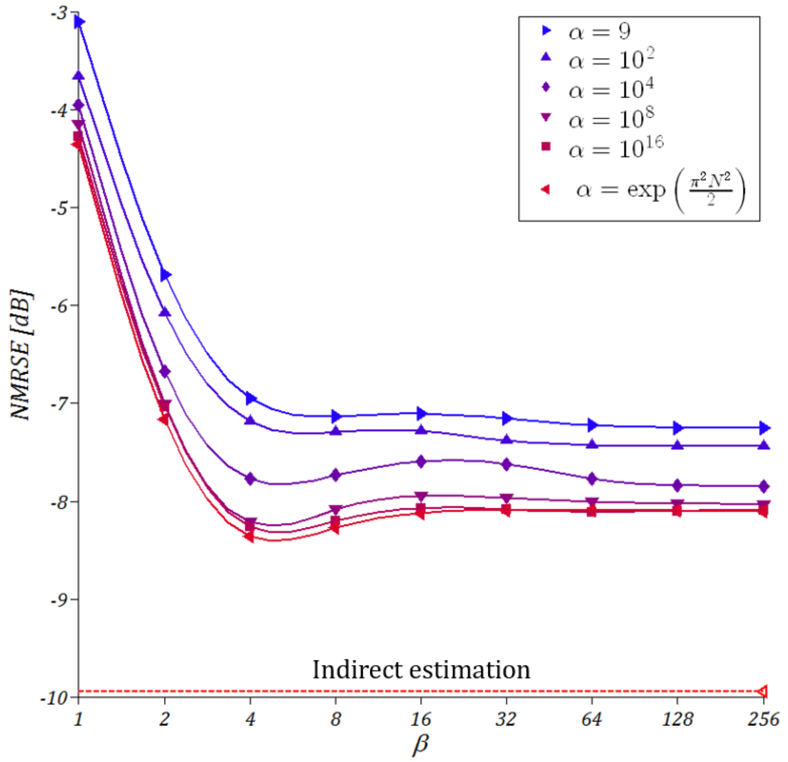

FIGURE 2.16: Reconstruction accuracy in NRMSE of the filtered signal directly synthesized through IF estimation using different pairs of α and β, in comparison with that of the indirect estimation from the peak-interpolation of the N^2-point ST.

where $y[n]$ is the signal synthesized from the filtered ST, and $\varepsilon[n]$ is the error defined as $x[n] - y[n]$. Fig. 2.15 depicts the reconstruction accuracies achieved with the octave reverse, and the TI, implemented using different combinations of α and β. The reference is the NRMSE calculated between $x[n]$ and $s[n]$ as in Eq. 2.57. We observe that the NRMSEs of the TI are not strictly monotonic functions with respect to β, since there are a few inflections points barely noticeable (i.e., by fractions of decibel). Their presence is not meaningful because it is justified by two reasons. First, not every sampling scheme contains the TF components sampled by lower-density schemes, since the frequency domain is sampled nonlinearly. For instance, given constant α and increasing β, the new and larger set of frequency indices does not necessarily include the old and smaller set that was sampled with a higher spectral compression ratio. This fact follows from the nonlinear growth defined in (**??**) and is evident from Fig. 2.9. The second reason is simply due to the error introduced by interpolation. As for the FI, the resultant NRMSEs in Fig. 2.15 are not visible because they are higher than the reference error for any α, except when $\alpha = \exp\left(\pi^2 N^2/2\right)$, which returns -10.94 dB. These results confirm that the FI is only suitable to retrieve the filtered signal from the fully redundant ST, and is not capable to cope neither with time decimation nor spectral compression. Alternatively, the synthesis can be carried out through IF estimation. The results in Fig. 2.16 are obtained using Eq. 2.29 to directly estimate the amplitude and the phase of the filtered signal. Instead, the indirect reconstruction based on IF peak-interpolation estimation was found to perform well only without compression; therefore, for the sake of comparison, the respective NRMSE is calculated only for the original N^2-point ST. An important observation can be made from Fig. 2.16 that is increasing β above a certain value does not bring significant improvement to the signal synthesized through IF estimation. Overall, the results suggest that the applications that employ the ST to modify the analyzed signal TF characteristics can benefit in terms of accuracy from a flexible TF sampling scheme. In fact, a level of redundancy aids the discrimination of the TF components to be filtered, and it also lessens the artifacts inevitably distorting the information due the filter mask itself. For example, Fig. 2.15 shows that the NRMSE falls by more than 3 dB with respect to that obtained using the FST when β is increased from 1 to 2, while the size of the corresponding ST is still less than $2N$, as clear from Fig. 2.10. Finally, the pair of parameters α and β can be tuned to rescale the TF representation according to a controllable trade-off between complexity and reconstruction accuracy.

2.7 Conclusions

After addressing in depth several aspects of the discrete ST, a complexity-scalable version of this transform is built upon a generalized time-frequency sampling scheme. The scheme is specifically devised to grant more flexibility for fast implementations. Our scheme can arbitrarily sample the two-dimensional representation in the TF domain in order to compromise between practical computational burden for analysis and decent reconstruction accuracy after synthesis. A controllable amount of redundancy allows for extensive modifications of the TF components, which instead implies significant errors if the existing and non-redundant schemes of the FST and the DOST. In this framework, we have assembled the forward and the inverse transforms into a digital system that can analyze any input non-stationary signal through a multi-rate bank of lowpass Gaussian filters, and synthesize an output through an equalized filter-bank summation. The rates and the sub-bands of the channels are designed according to the chosen sampling scheme and their parallelism is suitable to fulfill real-time constraints. The complexity-scalable architecture proposed ultimately enables advanced processing techniques for filtering a signal through the complexity-scalable ST, which can be useful in diverse and novel applications. A practical example is the rejection of in-band and pulsed disturbances in a specific case study. The formulation provided may be also extended to data structures of more than one dimension.

Chapter 3

Interference Source Localization

Critical infrastructures relying on communication and navigation systems (e.g., public health and transportation, telecommunications, and security) have been recently experiencing disruptions due to outages caused by interference events with increasing frequency. When intentional, these attacks are carried out by mean of devices named jammers, which deliberately broadcast interference. In order to quickly restore the full operation of these services, finding the sources of interference is a necessary task for taking actions against their threat, especially when the origin is malicious. This demand has motivated the research of methods and the development of systems for identifying and localizing terrestrial transmitters of unknown features in diverse fields of electronics engineering. What follows is a comprehensive review of the state of the art regarding interference localization and similar general-purpose applications. When it comes to the GNSSs, the aim of this extensive review made is to highlight two aspects of the problem of interest.

1. Most of the techniques in the literature tackle the presence of only one emitter at a time. However, the proliferation of (in-car) jammers among civilian users has recently made more and more realistic the possibility for a receiver of undergoing jamming attempts from *multiple* sources at different locations. This unfortunate but likely situation is an issue of growing concern in the GNSS community and is still open, despite the plenty of solutions theorized and implemented in the last decades.

2. The vast majority of jammers on the market nowadays transmit highly nonstationary signals in the form of sawtooth waveforms, which may be regarded as *signatures* of their spectral characteristics over time. Novel localization systems should take advantage of this fact, when retrieving the observables necessary to infer the locations of the devices responsible of interference. This aspect opens up to the utility of TF analysis for jammer localization.

In this context, the present dissertation has the merit of applying the desirable properties of the ST to the localization of simultaneous jammers, for the first time to the best knowledge of the author. The original contributions may be summarized by the following list according to the order of appearance in this chapter:

- the in-depth evaluation of the tracking performance for the position and velocity of a single jammer with time and frequency measurements, focusing the investigation on the impact of the signal snapshots, which has been usually overlooked in similar works;

- the innovative use of the ST for identifying the number and the characteristics of multiple jammers by exploiting their inner periodicities, hence overcoming the limitations of the state-of-the-art and STFT-based characterization techniques for single jamming waveforms;

- the combination of the ST-enabled jamming characterization together with the cooperation of manifold receivers for tracking simultaneous jammers of different characteristics and power levels.

Therefore, the main contribution is a novel and promising algorithm for recognizing and separately localizing simultaneous jamming waveforms. The adoption of the ST in this application enables a step forward in terms of capabilities, by shifting the resolution of the estimation ambiguities and the data associations from the conventional delay-Doppler domain to the more convenient TF domain. The complexity issue addressed in Chapter 2 does not threaten the feasibility of the approach proposed, because the algorithm is supposed to run on a back-end server, rather than being distributed on the individual receivers. A proof-of-concept case study was simulated to show the numeric accuracy achievable in ideal conditions of line of sight. The downside of this approach is that jammers are distinguishable only when they feature different TF characteristics in terms of periodicity. Nonetheless, as explained in the last section, such a shortcoming could be overcome by exploiting also the difference in received powers and possibly in the angles of arrival.

3.1 State of the Art

Given the extension of the scientific literature on this topic, we may narrow down the state of the art to the **cooperative** and **passive** localization of **terrestrial** and **mobile** transmitters with **little or no knowledge** about their signals. The rationale behind this choice comes from two facts. The sources of interference have usually *unknown* number and characteristics and, for practical reasons, their locations can be inferred only by means of *indirect* and *repetitive* observations. In other words, we focus on localization methods that cope well with interfering transmissions by collecting measurements related to their origin through a physical or synthetic set of *arbitrarily* and *spatially* distributed sensors. Hence, they need neither any active and direct interaction with the interferers (e.g., round-trip time estimation with radar pulses or data frames) nor any knowledge about the waveform transmitted or the propagation environment (e.g., for multipath exploitation). Despite the specificity of this context, such an application is shared by countless research papers in different branches of electronics engineering. Particularly, we may distinguish among methods for localizing generic radio-frequency (RF) emitters, signal and interference sources within wireless sensor networks (WSN), and jammers for GNSS receivers. These three fields of application are one by one summarized in the following subsections.

3.1.1 Generic RF Emitters

Most of the state-of-the-art systems aim at localizing a ground source from the respective RF emission in a general framework, which is suitable to a plethora of applications, including radar, sonar, acoustics, wireless communications, positioning, and navigation. To categorize the methods complying with the application of interest, first of all, we may discriminate between the methods for a single emitter localization based on trilateration (or multilateration) and triangulation, with a final paragraph on the systems capable of tracking multiple targets at the same time.

The localization methods reviewed in this section are implicitly based on a *centralized* and *two-step* approach that is also adopted hereinafter: from the signals received at different sites a central processor retrieves the observables, which are then used used for computing the location estimate. These methods are, however, inherently *suboptimal*, because the measurements are processed independently, although they are tied to the same emitter. The resultant estimation is so not guaranteed to yield the optimal solution. Nevertheless, it still approaches the Cramér-Rao lower bound (CRB) when the observables are *asymptotically* and *mutually* uncorrelated under mild conditions, which include moderate signal-to-noise ratios (SNR) and a large number of observations. In this case, the two-step approach is equivalent to a single one, whereas an optimal direct estimation process is useful only for low SNR and/or short signal samples, as demonstrated in [49].

Trilateration

Under this category fall the methods that infer the relative and apparent point-to-point distances, namely the *pseudoranges*, of sensors from an emitter, with received signal strengths and time differences of arrival.
The observation of the *absolute* time of arrival (TOA) of the signals coming from several satellites underlies the self-localization capability of GNSS receivers. By comparison, the application of interest may be regarded as an inverse operation, which is further complicated by the lack of knowledge of the emission time. This unknown can be eliminated by observing the *relative* delay of the signal received by two synchronous sensors at known locations, at the expense of a strengthened noise intensity and correlation in the measurements. Generally, differential timing information may be obtained through baseline interferometry, time difference of arrival (TDOA) estimation, spatial spectra, phase antenna, etc. The advantage of TDOAs is that they can be observed by digitally processing the time-aligned signals that are captured by different front ends, without requiring an antenna array or any other analog component. A sensor-side synchronization mechanism relies on the GNSS time reference in the most straightforward implementation. One of these observables restricts the possible locations of the emitter to a hyperboloid having the two sensors as foci. Thence, in a two-dimensional scenario, hyperbolic location estimation pinpoints the signal source at the intersection of the two or more hyperboloids that correspond to the TDOAs from three or more sensors. Ideally, when there are no errors on the measurements and the sensor locations, the emitter would be exactly localized. However, the intersection of hyperbolic curves does not determine a point but an uncertainty region in space, due to the inevitable noise and imprecisions. Since the hyperbolic equations relating the TDOAs to the respective differences of pseudorange, the estimation of the unknown emitter location is a highly nonlinear optimization problem and, more specifically, one with a non-convex function. If a fully *efficient* (i.e., minimum variance and unbiased) estimator exists, that is the maximum likelihood (ML) one. The ML estimator *asymptotically* produces an *optimal* solution to this problem, attaining the CRB. This is true as long as there is a large number of mutually independent observations to increase the immunity to additive and non-systematic measurement errors. Another important property is that this solution is statistically equivalent to the one of a nonlinear least squares (LS) estimator, if the zero-mean and white measurement noise has a Gaussian probability density function (PDF). This second estimate is deterministic rather than statistical, as it is obtained by minimizing the sum of the square residuals, without so any probabilistic assumption on the observation likelihood. To avoid the hassle of an exhaustive numerical search in the solution space, the computational expense of the resolution is usually reduced with approximations of the ML estimator, which are nearly optimal when specific simplifying conditions are met. On the one hand, a closed-form fix can

be formulated by rearranging the nonlinear equations into a set of linear ones with the introduction of extra variables. On the other hand, these nonlinear equations can be linearized via Taylor-series expansions. Both these strategies are ultimately employ LS estimators and may be found in [50, 51] and in the references therein. Particularly, Taylor's theorem is adopted by the Gauss-Newton algorithm to give a local linear LS solution around an a-priori estimate to a set of nonlinear equations. These equations generally suit any localization problem, regardless of the measurement nature, as firstly emphasized in [52]. This is a classic method for iteratively improving the estimation accuracy, starting from an initial guess. The global convergence is not assured and depends on the a-priori estimate, which should be close enough to the actual solution to avoid the local minima of the function. In the following sections, we resort to the same *linearization* albeit in a *recursive* fashion instead of *iterative*, because it leads to estimators mathematically simple, commonly used, and compatible with any number of sensors.

The relative motion between the emitter and the sensors cause the signals to be received with Doppler frequency shifts. The observation of this phenomenon can be exploited for localizing a source, given that all the clocks used for sensing are consistent in both time and frequency. Indeed, one frequency difference of arrival (FDOA) defines a surface of possible emitter locations, the shape of which does not depend on the carrier frequency, but on the sensor geometry and relative velocity. This principle is used for localizing both a *stationary* emitter with sensors maneuvering along known trajectories and a *movable* emitter with sensors fixed at known locations, such as in [53] and [54], respectively. The combination of TDOA and FDOA measurements enables the extension of these methods to a generic scenario with moving transmitters and receivers. The incorporation of Doppler frequency shifts is also useful for reducing the number of sensors required, enhancing the location accuracy, and estimating the emitter velocity. As a consequence, the extension of the resolution strategies aforementioned to the now two-dimensional estimation problem has been widely investigated in [49, 55–59]. Innovative ways to find the global solution of the non-convex optimization function make use of a Monte Carlo importance sampling method throughout the scientific literature in [60], apply semidefinite convex relaxation in [61], and multidimensional scaling analysis in [62]. They perform better than conventional strategies based on LS estimation even with large measurement noise, but are costly in terms of complexity.

Despite the fact that measuring frequency differences between received signals requires sufficiently accurate and stable local oscillators mounted on the various sensors, a recent research trend has investigated the use of this combination for location estimation with basic sensing platforms, like unmanned aerial vehicles (UAV). For practical reasons, much effort was devoted in [63, 64] to use only two sensors that outmaneuvers the emitter to be localized. Initially relying on solely TDOA observables and lately switched to TDOA and FDOA measurements, the methods proposed make use of the observation time history to avoid the need for more physical sensors. In this regard, the aforementioned LS estimation can be employed to process the sequence of observables in *batches*, given a deterministic model of the emitter trajectory. However, this iterative process is computationally intensive, because it is repeated whenever a sufficient number of new data are available. In addition, the previous datasets are stored to be at disposal. Since two sensors can only provide one TDOA or one TDOA/FDOA pair at time, besides the processing time, the wait between two consecutive estimates might be excessive, especially in dynamic scenarios. For a matter of execution time and efficiency, modern localization systems

then prefer recursive estimators, which avoid to reprocess all the observations without compromising on the final accuracy. More specifically, the state of the art mostly refers to a variety of Kalman filters (KF), which *sequentially* update the current estimate upon the arrival of new measurements and which rely on a statistical model of the emitter motion. As opposed to the simpler LS, this choice entails the adoption of a Bayesian philosophy, in which the goal is the minimization of the MSE. An historical perspective on this topic is given in [65]
The TDOAs can be derived from the one-dimensional cross-correlation function of the raw signal samples received from each pair of sensors. The joint estimation of TDOA/FDOA pairs is based on a cross-ambiguity function (CAF) that essentially extends the cross-correlation to account for the differential frequency shift, other than the delay. This two-dimensional function is so defined over a grid in delay-Doppler domain. A search for cross-correlation peaks through a ML estimator retrieves reliable observables, as long as there exists only one transmitter that has access to direct propagation paths to both the receivers. This approach is not effective in the presence of multiple emitters, the signals of which interact with one another, creating ambiguity in the cross-correlation function. Different strategies can resolve this uncertainty, as we will have the opportunity to see in more detail later in this chapter. Under non-line-of-sight and severe multipath conditions, instead, more complex techniques are necessary. This issue is beyond the scope of the present dissertation and has been widely addressed in the last decades by plenty of research papers, such as [66–68], just to name a few.
A simpler albeit less accurate alternative to TDOAs for *powerful* emitters and non-line-of-sight environments consists in extracting raging measurements from the received signal strengths (RSS). They are obtained without any synchronization mechanism and are usually readily available in low-complexity devices. From the knowledge of the exact channel path-loss model relating the power levels to the respective distances from the source, the two-dimensional location of this latter can be estimated with three or more sensors. This scheme, however, entails the access to trustworthy information about the antenna radiation patterns, the height of the transmitter/receiver pair, and the electromagnetic phenomena affecting the channel. For instance, the effects of shadowing and multipath-induced fading caused by obstructions, ground reflections, etc. of the surrounding environment should be evaluated. The difficulties in relying on RSS indicators of WiFi devices are experimentally confirmed in [69] with a two-ray path-loss model. Even when the assumption of omnidirectional antennas and the propagation model adopted match the reality, one should know either the transmission power spectrum or the power spectra received at reference locations within a region surveyed beforehand. Similarly to TOAs and TDOAs, the need for estimating the emission carrier frequency and power can be relaxed by resorting to differential RSS (DRSS) measurements. One of these observables defines a sphere of possible signal source locations in between the two sensors. This workaround, though, does not exclude the need for a path-loss survey. Nonetheless, despite the coarse accuracy they tend to provide compared to TDOAs, the signal strengths could be the only viable metric for location estimation with an infrastructure made by low-cost platforms. For this reason, they are popularly used in WSNs, which we address in the next subsection. A comparative evaluation of the performance of different LS estimators with DRSSs is reported in [70].

Triangulation

This category encompasses the direction-finding systems for estimating the location of an emitting source as the intersection of bearing lines, which originate either at multiple points along the known trajectory of a moving sensor or at multiple fixed locations of various sensors. In three dimensions, these lines do not necessarily intersect and LS estimators are usually put in place to resolve the spatial uncertainty, such as for TDOA hyperboloids. Discussions about manifold aspects and statistical solutions of this problem for a stationary emitter date back to the late sixties up to the nineties. They mainly come from aerospace engineering and may be found in [71–73].

The directions of bearing lines are obtained from angles of arrivals (AOA), which are measured by observing the beam pattern of an antenna array resulting from a *beamforming* process. This equipment is used either to electronically steer the beam of the radiation pattern towards the direction of highest signal strength or to measure the phase shifts at the individual antenna elements. Methods for location estimation based on AOAs only need two or more observations taken at different sites in a two-dimensional scenario and, more importantly, they do not require the time synchronization of sensors such as for measuring TDOAs and FDOAs. Another interesting advantage over the time observables comes from the fact that several angles can be measured from an individual device, which maneuvers to provide the spatial diversity of a synthetic array of sensors. Consequently, there is no need for data links to transfer raw signal samples between the sensors and the processing unit in charge of the two-step estimation process. The processor can be hosted on the sensing platform. Therefore, systems based on AOA are usually simple and able to operate autonomously as well as covertly. Compared to time measurements, angle are however more sensitive to errors as the range between transmitter and receiver increases and similarly susceptible to multipath propagation, which heavily degrades their accuracy indoors and in urban settings. Furthermore, the size and the calibration of antenna arrays might be unsuitable to inexpensive devices of opportunity (DoO) and light vehicles (e.g., UAVs).

For a single and movable platform, Doppler frequency shifts represent a natural add-on to bearing measurements, as firstly studied in [73]. A publication of the same author in [74] cites more recent methods for location estimation and also provides references to studies on the target observability with respect to the sensor manoeuvre. Conversely, when it comes to multi-platform sensors, the works published in [75, 76] propose to combine TDOAs with AOAs, by taking advantage of the complementary availability of these two observables. With the same measurements, the decentralized location estimation capable of working without line of sight is introduced in [68] through the expectation-maximization algorithm, where the computation of the statistics is distributed over the sensors communicating with each ohter. A special mention finally goes to the single-sensor method presented in [77], which is fully *hybrid*: it initializes an iterative LS estimation with a coarse search over the three-dimensional space of angle, heading, and frequency measurements for different possible emitter altitudes. These observables are jointly estimated with parallel CAFs, which are computed between the direct-path signal and the replicas reflected by the terrain at different angles. Thus, the spatial diversity necessary to resolve the nonlinear equations is obtained through the scattering environment, without any prior knowledge about the locations of the virtual sensors. This remarkable lack of assumptions distinguishes this system from most of the later methods exploiting multipath for localization, which assume known scatterers (i.e., a database of ray-tracing data) and, therefore, do not comply with the application of interest.An overview on this topic is available in [78].

Tracking of Multiple Targets

Without making simplifying assumptions, the passive localization of an unknown number of emitters is clearly a more complicated task compared to the same application for a single emitter. This challenge has been historically tackled by using the signal received at the elements of one antenna array. The well-known paper in [79] is arguably the first one to introduce the use of subspace algorithms for localizing narrowband sources. It estimates the AOAs of the emitters by analyzing the eigenvectors of a cross-power spectral density covariance matrix with the multiple signal classification (MUSIC), which enables the distinction of frequency components due to different emitters. Alternatively, the estimation of signal parameters via rotational invariance technique (ESPRIT) relaxes the requirements in terms of computational power and precision of the array calibration at the cost of more antenna elements. Following the emergence of multiple-input multiple-output (MIMO) systems, these algorithms have become of renewed interest for wireless communications and bistatic radars, for instance in [80, 81], with efforts to reduce their often prohibitive complexity. With MUSIC and ESPRIT, a N-element antenna array can resolve up to N-1 sources, whereas other techniques are necessary to increase the degrees of freedom (e.g., [82]). Instead, when a two-dimensional scenario is a reasonable approximation of the reality, the simplest possible way of localizing multiple emitters based on AOAs is clustering the intersections of redundant bearing lines, which depart from a movable platform or more than one. When the use of possibly cumbersome antenna arrays is not practical, another option is to resort to the TDOAs that can be simply obtained by digitally processing the signals from sensors synchronized in time. In such a multi-platform arrangement, the challenge is moved toward the *data association* problem: the set of correlated measurements collected by the sensors should be grouped into subsets, each of which is originated from the same target for a number of emitters not known a priori. A solution based on the assignment formulation is proposed in [83], which adapts the multiple KFs tracking the signal sources and recursively exploits their prediction estimates to save computations. By extension, the same problem affects TDOA/FDOA pairs. A similar spatial uncertainty is also present in bearings-only systems, when the AOAs are measured by asynchronous sensors deployed at distance from each other. In this situation, the algorithms described in [84] recursively estimate the target existence probability, by also taking into account clutter measurements that randomly arrive at the antenna arrays. It is worth mentioning that this problem has been extensively addressed since a while for localizing sources of acoustic emissions with TDOAs, by means of more advanced techniques (e.g., expectation maximization), such as those cited in [85].

Other methods for multiple emitter locations in [86] circumvent the data association problem by considering *disjoint* emitters. The disjointness property states that the signal sources do not interference with one another either in time of frequency. In other words, they are confined over separate time interval or frequency bands, which make their respective sets of TDOAs and FDOAs distinguishable. This is a very convenient property that does not realistically apply to the subject of this thesis.

Particle filters (PF) and, more generally, sequential Monte Carlo methods provide a powerful and flexible *heuristic-like* tool for recursive estimation. Over the last twenty years, their ability to simultaneously track posterior probability distributions have progressively being applied in various ways to solve both the two distinct problems of data association and multiple target localization at once, as examined in [87] and later works. This happened as well for many other applications, such as positioning (e.g., [88]). A premise is necessary to understand the potential of PFs. The KF is a well-established minimum MSE estimator that is optimal, and so equivalent to the fully-efficient ML estimator, if the joint PDF of the state (i.e., the unknowns) and the measurements is Gaussian. Necessary but not sufficient conditions for this property to hold are the linearity of the stochastic models for both the state prediction (e.g., the emitter motion) and the observation, and the Gaussian nature of their respective noise distributions. This is clearly not the case for the highly-nonlinear class of estimation problems, of which localization belongs to, especially under non-Gaussian noise. Besides suboptimal versions of KFs, another solution is, therefore, resorting to PFs, which are not subject to the restrictions imposed by the linearity and Gaussianity hypotheses. The basic idea behind them is the exploration of the state space with a grid of random probability masses, also known as particles. In brief, a weighted set of particles approximates the PDF of the state conditioned to the observations through importance sampling. At every recursion, this set is propagated through the prediction, which is performed for every particle separately, and their weights are updated according to the likelihood of the new measurement combined with the former weights. In between these two steps, a *resampling* operation is in charge of dismissing the particles with lower weights in order to concentrate them where the likelihood is supposedly higher. It is relevant to the discussion to underline that several issues affect the behavior of these filters, namely the poor accuracy in high-dimensional state spaces, the risk of particle depletion, the difficulty of evaluating the necessary amount of computations in advance, and the significant complexity compared to an equivalent array of KFs, just to name the main ones. These aspects are indeed subjects of an active area of research that is out of the scope of this state-of-the-art review. An overview on Bayesian filtering is given in [89]. An example of PF developed for multiple target localization is in [90]. Plenty of others may be found in the literature, especially for acoustic sources and WSNs. Mixed implementations of KFs and PFs also exist. The method described at the end of the present chapter for a similar application scales an array of multiple single-emitter trackers (e.g., KFs), rather than making use of one multi-emitter tracker with a PF.

3.1.2 WSN Non-Cooperative Sources and Jammers

The general framework concerning the passive location estimation of *non-cooperative* signal sources has been ported to the WSNs of anchor nodes, in the direction of higher efficiency in terms of energy and communication bandwidth. In this context, examples of systems enabled by TDOA and/or AOA measurements are in [75, 91]. Instead, the supposed defect of collaboration excludes all the methods based on absolute TOAs, because they presume the establishment of a source-node message protocol for synchronization. The observation of TDOAs is achievable by equipping the sensors with GNSS receivers or low-drift clocks accurate at the nanosecond level (e.g., atomic clocks). Since these implementations are hardly viable for WSN, the publication in [92] presents an ML estimator for passively estimating both the location of multiple emitters and their internal clock offsets with asynchronous nodes. For the same reason, networks of low-power and low-cost devices oriented to the Internet of things (IoT) prefer less sophisticated techniques, either distributed or centralized, which utilize RSS measurements. This choice is made in manifold scientific papers related to this field, for instance in [93, 94] and those cited in there. However, they usually assume the a-priori knowledge of the source transmit power. Locations and powers of multiple sources are instead jointly estimated through a sequential Monte Carlo method in [95]. Theoretical and experimental comparisons between the location accuracies achieved with TOAs and RSSs are provided in [96] and [97], respectively.

A whole branch of communications engineering research is dedicated to the denial-of-service attacks carried out by wireless network jamming. There is indeed a broad variety of interference attempts from the physical to the network layer, and, consequently, of countermeasures, which are surveyed in [98, 99]. Jammer localization methods mostly leverage the effects of interference on either the network topology or the power levels received at the various nodes. The first approach relies on *range-free* metrics, such as packet delivery ratios in [100], lists of neighbor nodes and hearing ranges [101], and changes and geometries in the connectivity matrix of nodes [102–105]. These metrics are sensitive to the node density and distribution in space as well as variations and irregularities in the propagation environment. Even under favorable conditions, they suffer from poor resolution and reliability. As a consequence, they are effective in localizing from one jammer to a *known* number of jammers, though only when the interferes are either far from each other or sequentially turned on. The second approach offers potentially higher accuracies by inferring the multiple jammer locations through the respective signal strengths. Their transmit powers are in turn estimated from the packets collected by the jammed nodes. However, the estimation of the jamming signal strengths is a challenging task, because interference is embedded into the regular network traffic. This challenge is overcome in [106] for the simple case of jammers that constantly transmit regardless of whether the channel link is busy or idle.

The difficulties in retrieving AOAs, TOAs, and TDOAs in persistent non-line-of-sight conditions has motivated research on techniques that create a database of unique *fingerprints*, namely received signal features, with respect to the possible source locations. Conventional fingerprinting relies on pre-trained and predicted maps of received power levels or certain information about multipath-rich channels. The measurements collected are then matched with these maps both for positioning and localization in WSNs (e.g., WiFi and WiMAX), which are deployed either indoors or in dense urban settings. Works based on RSSs in this field may be found in [107, 108] and the citations therein. However, the matching algorithms are effective only within the well-surveyed regions of the propagation environment. Moreover, both the complexity and the lack of knowledge of the source signal center frequency, structure, and time-varying transmit power complicate the feasibility of fingerprints based on RSS and channel impulse response. All these aspects are highlighted by a recent publication in [109] that proposes the use of AOAs of the scattered signals as fingerprints, which are extracted from the phase difference between the elements of an antenna array, weighted according to a cross-correlation coefficient between antennas, and interpolated in the spectral and spatial domains.

3.1.3 GNSS Receiver Jammers

The intent behind the previous subsections is to widen the discussion on passive localization systems of interferers “beyond GNSS”. Indeed, it is the opinion of the author that the field of existing jammer localization methods in the context of GNSS receivers is still not as mature as adjacent research areas, which deal with the localization of generic emitters, especially when it comes to multiple targets.

An up-to-date review of this topic in the field of countermeasures to interference

for GNSSs is available in [110], which we retrace to complement the references of this subsection, lingering on the most relevant papers cited by it. Particularly, we focus on cost-effective *software* techniques based on trilateration, which is indeed suitable to the use of inexpensive and portable DoOs. Such a restriction excludes triangulation, because the direction-finding capability is enabled by additional bulky *hardware* at the receiver front end. Interference localization methods based on AOAs have been investigated since the early works concerning GPS, for instance in [111, 112]. Predictably, they proved to achieve average accuracies highly dependent on the jammer-sensor distances and the measurement errors, up to hundreds of meters. Nowadays, newer systems tend to partially or totally abandon direction-finding techniques, because bearing measurements are not suitable to dense urban scenarios, which is where the highest concentration of jammers is likely to occur.
Once they arrive to the receiver, GNSS signals reside below the noise floor and are so easily overpowered by any interference transmitted nearby. Nevertheless, this vulnerability has a bright side: any anomaly in the RSS could reveal jamming attempts and indirectly localize the respective origin. Standard GNSS receivers are equipped with an automatic gain control (AGC) that is responsible for minimizing the quantization loss by adjusting the gain at the output of the low-noise amplifier (LNA) in order to spread out the input signal over the full dynamic range of the following analog-to-digital converter (ADC). The larger is the finite set of bits, the less sensitive is the SNR degradation to the variable gain of this second amplification stage. The AGC/ADC loop is tuned by a feedback voltage, which can be exploited for measuring the power received from strong signals. This in-built utility avoids the need for special-purpose implementations and is the reason why jammer localization systems have been developed based on RSS observables acquired from the AGC, such as in [113]. In reality, though, the domain of the nonlinear voltage function with respect to the emitter power and distance is actually restricted within a two-dimensional region due to the front-end limitations. The higher bound is caused by the AGC/ADC saturation, which likely occurs if the receiver is in proximity to the jammer, unless a switchable attenuator is included. The higher bound is because the voltage resolution may be insufficient for low-power or distant emissions. Such limitations imply that the useful measurements can be collected only from the receivers not too close to or far from the jammer. The use of the AGC is also suggested in a paper in [114], which outlines the main challenges in creating a national infrastructure to detect and localize jammers for both GNSS receivers and cellulars. This system is built around the *crowdsourcing* of AGC voltages from a *multitude* of DoOs, which are usually mobile phones. Every phone at a suitable distance from the jammer could report to a central server its own position and orientation (i.e., the antenna gain) together with a jammer-to-noise power ratio (JNR) estimate. This measurement is obtained by a meter that observes the RSS from the non-saturated AGC. The poor accuracy of the individual observations is compensated on the server side by aggregating the measurements from a huge crowd of devices. And even so, measurement noise, the propagation effects and the uncertainty on locations of the jammed devices may lead to significant biases, depending on the scenario. For instance, according to the best-case scenario described in [114], less than 30 meters of accuracy is achievable by processing the LS best fit over a grid of hypothetical locations for the JNRs reported by a thousand phones, which uniformly cover a 1-km square area. The exchange of AGC levels through the infrastructure of vehicular ad-hoc networks is also being mentioned in [115] as a countermeasure to in-car jammers, which could compromise the safety of future intelligent transportation systems. Likewise the AGC voltage, a metric readily available from the components off-the-shelf inside GNSS receivers

is the *post-despreading* C/N_0 estimate for every tracking loop. Hence, drops in the satellite signal strengths represent a basic tool for detecting and localizing jammers. Recently, the paper in [116] explored the accuracy of a stationary jammer location estimation with field observations of C/N_0 levels, that are collected around the interferer both from a two-row grid of 15 mobile phones and a sensor synthetic array made by one moving device. The experiments carried out outdoors show errors lower than 10 m, which nonetheless result from a crowd area or a sensor trajectory that lie within few tens of meters. Moreover, these tests neglect the realistic and adverse conditions that put in question the feasibility of any localization method solely relying on C/N_0s for real-world applications. The most important shortcomings include the lack of information about the in-built C/N_0 estimation of GNSS chipsets (e.g., examines manifold choices), the possible poor sensor geometry, large distances, and environmental small- and large-scale variations. The same crowdsourcing approach is also studied in [117], where an a-posteriori probability map of a single jammer location is built through a PF. The interferer is localized at the intersections of probabilistic coverage circles, which are simply modeled with free-space attenuation. Again, both in [116] and [117], the tacit supposition is to have a huge redundancy of measurements that compensates for the discrepancies between the model and the reality. As an alternative for large outdoor environments, where adopting a path-loss model that mirrors the actual behavior of the propagation channel is possible without surveying, a generic localization method is studied in [118]. Here, a GNSS interference scenario is replicated with WiFi devices that collect DRSS measurements. The range of radial accuracies obtained is on the order of tens of meters with four or five sensors within an area less than 100 meters wide. The same scheme is adopted in [119], which performs a simulation campaign to assess the impact of several factors on the performance of a KF. Different numbers of randomly-placed sensors, log-normal shadowing variances, and quantization bits lead to location biases lower than 20 meters over a flat surface of 500-m side, when considering more than 100 sensors.

The reliability of power measurements tends to degrade for long distance and sparse sensor arrangements. The superior accuracy granted by TDOAs over RSS and DRSS observables is the reason why there exists a rich literature about jammer localization based on differential delays, which are measured between permanent or deployable and time-synchronized sensors. These latter ones may be distant from each other and need to be connected to a central unit that processes raw *pre-despreading* samples to retrieve the TDOAs, usually through cross-correlation. First efforts in this direction are the hyperbolic localization methods described in [111, 113]. Another example is the GNSS Environmental Monitoring System (GEMS), which was followed by various publications. The paper in [120] compares with numerical simulations the performance of several cross-correlation methods for TDOA estimation, which overall can achieve sub-meter errors in localizing a jammer even for negative JNRs. The scenario simulated assumes a free-space propagation between three equally-spaced sensors delimiting a 4km-square area and the interference source located in the line of sight at the center. The incorporation into the sensors of GEMS of direction-finding capabilities is demonstrated in [121] to bring several benefits to the TDOA accuracy, by controlling the receiver radiation pattern in a twofold way. At first, the antenna array can steer a null at the jammer and a beam at one of the visible satellites in order to enable precise synchronization in time between sensors in spite of the in-band interference. Then, once the AOA of the jamming source is established, steering a beam toward this direction improves the JNR and so the precision of the TDOA observation. The system sensitivity is thus potentially extended to intercept weak jammers and, equivalently, to cover distant ones. The analysis of aspects of TDOA estimation and the AOA addition is finally merged in [122]. The evaluation of stand-alone TDOA as well as hybrid TDOA/AOA localization performance of UAVs is inspected in [123] and [124], respectively. In the latter paper, the direction-finding capability is also used to aid the initialization and convergence of a KF. Although the great potential in terms granted by adding bearing measurements, a well-known issue in finding the sight lines of very weak jammers (i.e., -20 dB of JNR) is that GNSS signals are usually stronger than interference and thus they impair the AOA estimation. This is akin to the classical *near-far* problem, where a strong signal masks a much weaker one. For instance, this effect cause self-interference among satellites visible with far different C/N_0 levels. Existing solutions for preventing this issue inside GNSS receivers are summarized in [125]. Among them, a subspace projection technique is ported in [126] to the realm of jammer localization with antenna arrays. Similarly, the same technique is applied in [127] to reject GNSS signals, which create cross-correlation peaks that can be mistaken for weak terrestrial sources, thus limiting the sensitivity and so the coverage of many interference localization systems based on trilateration. Besides, these peaks could also interact with those produced by interference, thus causing a bias in the resultant TDOA estimates. A simple non-coherent cancellation technique is presented in [128], which essentially proses to subtract the mean absolute value of the non-interfered cross-correlated function from the interfered one. The publication in [129] gives an insight into the detrimental impact of the cross-correlation products of GNSS signals along with jamming ones on the detection and localization performance of the GEMS, depending on the distances among sensors. Briefly, if a satellite signal is received at two close locations, at exactly the same time and with the same sampling frequency, then a sharper cross-correlation peaks arises around zero as a result of the alignment in time. Since the same likely happens for the other visible satellites, their peaks sum up into one with significant amplitude, which can exceed that of the product of interference. As a consequence, this scenario is more prone to the aforementioned near-far issue. On the contrary, when the distance between sensors increase, the satellite signals are received at slightly different time instants, thus separating and smoothing their respective peaks, and the interference is easier to be discriminated. Another aspect of the GEMS, as well as any other system relying on TDOAs, is to ensure the synchronization of spatially-distributed sensors under jamming attacks. Indeed, any time misalignment between the signals received degrade the final location estimate. An experimental evaluation of the drifts experienced by low-cost and short-stability temperature-compensated crystal oscillators (TXCO) typically used in GPS receivers is carried out in [130]. These on-board clocks drive the down-converter to intermediate frequency or baseband and the ADC. As a consequence, frequency offsets cause both a spectral dilation and a shift of the signal in the sample record, which progressively contribute to the drift due to time offsets. Throughout the experiments, diverse jamming waveforms and transmit power levels are used to cause immediate loss of lock in all the channels of a receiver, which then goes into a holdover mode, namely using the local oscillator for open-loop synchronization. During the GPS outage, without being disciplined by the satellite clock, the responses of the timing and the frequency offsets do not show any noticeable drift at least before a certain number of seconds have elapsed after the interference is turned on. This interval of time is long enough to switch to a backup timing source or to perform a number of TDOA measurements. A longer holdover is possible with higher-quality oven-controlled crystal oscillators (OCXO) and atomic clocks, which are expensive and power-hungry though. Other options are the recovery of rough timing information from opportunistic signals, such as those of the base station of cellular networks and WiFi access points, and the exploitation of the multi-frequency (e.g., upper and lower L bands) and multi-constellation (e.g., GPS, Galileo, GLONASS) diversity. However, besides the cost that comes with it, such a redundancy is possibly useless against jammers that scan all the bandwidths allocated to these communication and navigation systems. This is actually the case of most civilian devices that are already commercialized online, according to [131]. For this reason, alternatively to the often impractical implementation of antenna arrays proposed in [121] at the sensor front ends, the work in [130] explores centralized post-processing techniques, which re-align the received signals in time and correct the TDOA estimates. The time and frequency offsets between the clocks of two sensors is determined offline, by demodulating the navigation data from a common-view satellite, namely visible at the same time at both sensors. The time-varying drift is modeled according to the intervals between correlation peaks of the same time-of-week message. The latency and the inability to acquire and track any of the GPS signals buried under the interference are the main concerns of using the post-processing techniques. Similar and yet simpler systems than the GEMS have been tested in [113, 132–134] by setting up a network of monitoring stations, which are equipped with low-cost USB front ends and connected to a server via TCP/IP. The server computer synchronizes the incoming data streams with the last time reference made available by a software GNSS receiver after the jammer is triggered. In this initial period, it is still possible to compute the exact time instants each sample was recorded together with the exact clock frequencies of the various front ends. Afterwards, the datasets are aligned in time and frequency by solely leveraging the TXCOs on board of the sensors. More specifically, a first-order model of the drift is used offline to correct the offsets by re-aligning and resampling the signals from different sensors before cross-correlation. The claimed response time to the jamming attack is on the order of few minutes, dependent on the data transfer speed and the computational power. The sensors are separated by large distances in order to avoid undesirable near-far effects and enhance both the coverage and the geometry. Precursor of the GEMS, the generalized interference detection and localization (GIDL) prototype described in [135] measure TDOAs by exclusively using closely-spaced and fixed baselines of antenna arrays, which are kept synchronous by sharing the clock from the a centralized processor via coaxial cables. This system is conceived as an upgrade to the Local Area Augmentation System (LAAS) for aircraft precision approach and landing. The extension of the interference localization capability to simultaneous jammers is investigated by the testbed developed in [136], which runs subspace (i.e., MUSIC) and LS fitting algorithms. For the sake of simplicity, perfect knowledge of the number of targets is assumed to be available from the analysis of the eigenvalues of the cross-power spectral density covariance matrix. This localization method draws from precedent seminal works about super-resolution techniques in [79, 137] for radar applications. The aforementioned issues of synchronization, near-field effect, and data association are addressed as summarized in the following. Sensors are equipped with the *tightly-coupled* architecture introduced in [138], so that they can lock at the nanosecond level to a common reference signal, which is coherently sampled together with the sample stream embedding the interference. The reference may be provided by signals of opportunity that, yet, may be themselves affected by the interference. As suggested by the authors, a simple and effective workaround is to pick up the signal of a GNSS satellite by pointing a directional antenna to the sky, which then would filter out any source at low elevation angles, including the jammers on ground. In such a case, because the reference lies within the same band of the interference, it is canceled before the parametric search for the TDOAs of multiple jammers with MUSIC and an iterative LS refinement. This subspace algorithm works properly under the assumption of interference flat power spectral densities, which holds whenever the jammed spectra are larger than the receiver bandwidth. Moreover, it is proven to have limited resolution for closely-spaced jammers, especially if their JNRs are relatively low and the integration time short. Once estimated, the TDOA measurements have to be associated to their respective jammers. The solution proposes a simple so-called phase closure statistic, which is computed by combining a triad of observables measured by distinct sensors and compared to a pre-fixed threshold to decide whether they come from the same origin. This principle is simple but can hardly be used to distinguish between the direct path and possible multipath replica. The field tests reported in the paper achieve a 20-m accuracy in an area less than 1 kilometer wide. Last but not least, the publication in [131] is meaningful for the purpose of the following discussion. Indeed, it proposes to exploit the typical *sawtooth* feature of GPS jamming waveforms in order to calculate the TDOAs in lieu of the usual cross-correlation. The aim is to reduce the communication bandwidth by shifting most of the computational load for the TDOA estimation from the central processor to the individual sensors. Every sensor transfers the estimate of its own time of arrival, which is equal to the interval between two consecutive crossings of the IF through a certain frequency. The IF is found at the peak of the power spectrum, which is computed with the FFT of sequential batches of samples. De facto, this procedure is equivalent to the IF estimation based on the STFT with non-overlapping windows. The significance threshold is set deterministically to some fixed fraction of the peak power in the current batch. Implicitly, the window size should be adapted according to the incoming waveform. Similar concepts have inspired the use of TF analysis for interference localization that is illustrated at the end of this chapter.

For the sake of comparison, the work in [139] analyzes the theoretical performance of any passive localization system, in terms of CRB and dilution of precision (GDOP), which are independently achievable by asymptotically-efficient ML estimators with TDOAs, AOAs, or DRSSs. The behavior of these systems change depending on the arrangement of the sensor array with respect to the jammer. Some of the consideration made are reflected by the results illustrated in the next sections.

Of particular interest are the developments made in [140, 141], the authors of which contributed to the research work described in the following. The former paper deals with the localization of closely-spaced jammers by taking advantage of the IoT paradigm and the potential of *cloud* computing. The application scenario outlined is that of a densely-populated areas covered by a variety of RoOs that synchronously gather *snapshots* of multiple jamming signals in a crowdsourcing fashion. The datasets collected are uploaded on the Internet to be processed by a cloud platform, which is supposed to have enough computational power at disposal to estimate and then cluster TDOA/FDOA pairs in nearly real time. The differential delay and Doppler frequency shifts are jointly estimated from the CAF. The ability to resolve the ambiguity in delay-Doppler domain due to the mutual interactions of more jammers is granted by leveraging the sawtooth pattern of jamming waveforms: the expected CAF is modeled as the output of a combination of input linear chirps scanning the whole receiver bandwidth. This model is then exploited to extract the TDOA/FDOA pairs as the iterative solutions of a minimization problem. Afterwards, a *clustering* operation sorts these observables in the position domain, instead of the delay-Doppler one: the measurement pairs are separated by comparing the respective time components to the differential delays that correspond to the points of a grid of possible positions. The clusters are formed wherever enough consensus

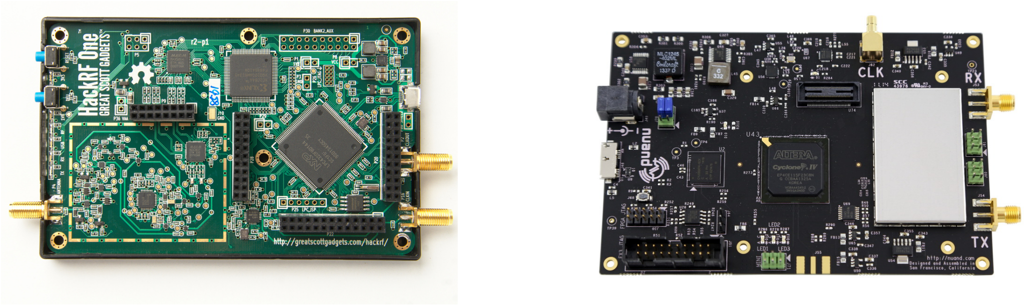

(A) HackRF One (greatscottgadgets.com/hackrf/).

(B) bladeRF (nuand.com).

FIGURE 3.1: Examples of front ends for DoOs.

is reached between the receivers, namely when enough observables (e.g., three) are close to the same point. This passage is essential to identify the number of jammers and overcome the issue of data association. Moreover, it allows for discarding the outliers due to false detections, measurement clutters, and model mismatches. For every cluster, the TDOA/FDOA pairs associated to it are fed to a weighted LS estimator that produces the source location. Given N-1 jammers, at least data from N devices are necessary to resolve the equations. A consistency check is made between the final solution and the initial position point. With four DoOs capturing 10-ms snapshots, the simulation results show accuracies on the order of few meters for two jammers lying at 40 meters of distance, inside the 100-m squared area between the receivers. The clustering approach seems a promising approach for tackling multiple jammers, despite their physical proximity and the effects of the surrounding urban or indoor environment. Alternative methods for multi-target tracking use super-resolution techniques with the same signal model of [140], since they likewise assume a flat spectrum over the visible frequency band. The paper in [141] examines an experimental setup for the synchronization of SDRs. These boards can record GNSS or LTE signals. The recording is triggered through a 1-Hz square pulse that is generated by either the internal clock or an external consumer-grade GNSS receiver. The mechanism described represents an enabling technology for the use of IoT sensors as DoOs in manifold applications, such as jammer localization and cooperative/opportunistic positioning. The experiments with the signal received from a single LTE base station confirmed that time offsets is practically bounded within one sampling period. The use of the on-board clock as trigger introduces a further frequency drift. Naturally, the same performance is achievable also with GNSS signals, whenever the number, the geometry, and the C/N$_0$ levels of the visible satellites are appropriate. The remarkable result is that SDRs operating at 20 Msps can be synchronized with offsets lower than 50 nanoseconds, when using GNSS time reference. This sampling rate is well within the reach of the specifications of popular and inexpensive boards like the HackRF One and the bladeRF, which are depicted in Fig. 3.1. Hereinafter, we will refer to similar front ends as the radios for low-cost and low-power DoOs.

3.2 Snapshot Tracking of a Single Jammer

In this section, we resume the application described in [140, 142] within a simplified scenario: a *passive* system for tracking the position and velocity components of a *single* and *movable* jammer is built around a cost-effective trilateration method, which is enabled by the *cooperation* of DoOs. Without loss of generality, the scenario is considered as two-dimensional. Therefore, we might define the *state* of the jammer as the column vector $\mathbf{x}_k = [x_k \, y_k \, v_{x_k} \, v_{y_k}]^T$, where (x_k, y_k) and (v_{x_k}, v_{y_k}) are the scalar components of the position and velocity, respectively, at the time epoch of integer index $k \in \mathbb{Z}$. Since four are the unknowns in this vector, then four is the minimum number of DoOs necessary to collect *simultaneous* and *independent* observables that are related to $\mathbf{x}_k$. These devices are *identical* transceivers equipped with basic front ends that *periodically* capture *snapshots* of raw samples, which are all *seamlessly* affected by the jamming attack. After reception, the sample records are sent to a central processor possibly hosted on the *cloud*. Here a digital and real-time estimation process takes place in two steps. At first, the snapshots coming from every pair of DoOs are cross-correlated to retrieve TDOA/FDOA pairs. Secondly, these measurements used to update the knowledge of a *recursive* and *suboptimal* nonlinear estimator, which tracks $\mathbf{x}_k$ based on a *linear approximation* of the observation equations. The same two-step approach is followed in [64], with the difference that here we take advantage of several snapshots over time by means of one conventional KF, instead of a more sophisticated bank of KFs. The recursion is also one of two major aspects that distinguishes this work from that of [140], where a weighted LS estimator is repeated independently for each sets of observables gathered at the current time epoch. The other difference is our restriction to a single jammer, which nonetheless we will drop later in this chapter. An overview of the application scenario described above is shown in Fig. 3.2.

The present subsection summarizes the findings published in [4]. The intent of this study is to get an insight into the impact of the snapshot *rate* on the final position and velocity accuracies. To focus the analysis on the effects of the snapshot density in time, some simplifying assumptions have been made to neglect other sources of error.

- The DoOs are *stationary* at known locations. – In reality, these devices might be moving and the inevitable inaccuracies on their locations increase the MSE according to the CRB derived in [55].

- The front ends are perfectly synchronous in both time and frequency. – This assumption is a usual simplification in the research literarute on jammer localization methods based on TOA, TDOA, and FDOA observables. It is here adopted to evaluate only the error due to the snapshot rate, without taking into account the error accumulating over time when the snapshots of different receivers cannot be aligned in frequency and, more importantly, in time. However, if we consider the DoOs initially synchronized to the nanosecond, for instance by using as time stamps the pulse per second provided by an onboard GNSS chipset [141], the impact of the synchronization error rising as soon as the GNSS signals are lost is actually negligible during an initial transient. Two facts help supporting this claim. First of all, since receivers fall back to the on-board clock in the absence of the GNSS time reference, a certain time alignment among them last for tens of seconds before the instability of common TCXO clocks (i.e., drifting up to 100 µs per day) could cause a meaningful degradation (e.g., beyond 50 ns) of the localization performance, as experimentally measured in [130, Fig. 5]. By that time, if the jammers causing the outage could be localized nearly in real time, their location are already known with high accuracy. Secondly, as argued in [136], a simple directional (e.g., helical) antenna pointing at the sky might enable the receivers to lock with one satellite, even under jamming attack. This is possible because the jammers are usually terrestrial sources at elevation angles comparable to the ground. Last but not least, any offset could be taken into account by simply increasing the noise affecting the measured TOAs or TDOAs, according to [141], without necessarily relaxing the assumption of ideally synchronous receivers.

- The jammer is continuously transmitting. – Procedures to reveal the undergoing attack and to handle the tracking accordingly are not necessary. Jamming detection is out of the scope of our analysis.

- The jammer is a *non-maneuvering* target moving at *constant speed* along a straight and stationary trajectory. – We may legitimize this model by supposing that the interval between two consecutive observations is short enough to average any slow acceleration in between with a local and constant velocity. A more general approach to track time-varying maneuvers is to ingrate the KF with an interactive multiple model (IMM) estimator, such as in [83].

- Both transmitting and receiving antennas are isotropic and the environment between them is fully characterized by a free-space path loss model equal in all directions. – This choice is a another common practice in the research literature of jammer localization methods relying on TOAs or TDOAs observables. For instance, it is made in [131] and many other papers cited throughout the manuscript, even though it contrasts with the reality of most of the application scenarios. Indeed, since the signal source to localize is terrestrial, the main obstacles in measuring the time of flight of the direct propagation path between transmitter and receiver are the lack of line of sight and the presence of multipath fading. The estimation of arrival times in a propagation environment affected by shadowing and scattering is by itself an open issue under research for many applications (e.g., positioning with signals of opportunity), which, however, is arguably beyond the scope of the topics in this thesis. Therefore, the choice made here is to overlook the generic and unfavourable propagation phenomena affecting the channel in favour of other and more specific aspects related to the signal processing. Of course, a final real-world implementation of the system shall take into account all these effects. Alternatively, the simplest workaround to avoid much hustle is optimizing both the arrangement and the radiation patterns of the receivers in order to avoid blockages and reflections by the surrounding environment, while still covering the area of interest. This is what usually done to install monitoring stations in outdoor environments, as for example around the airport in [143]. In the case of moving DoOs, however, the careful design of the receiver deployment is not feasible. Thence, besides sophisticated techniques (e.g., super-resolution resolution of multipath), a system of DoOs is supposed to count on a very large number of receivers and a massive availability of redundant observables, so that the measurements excessively corrupted (e.g., related to a delayed replica of the direct-path signal) can be discarded without interrupting the localization of the source. The principle of redundancy is exploited by crowd-sourcing signal records in [116].

- The interference power received by every DoO is always below the *saturation* threshold – The jamming attack has to come from a transmitter far enough from the receivers not to saturate the front ends. In the close proximity, in fact, the LNA is forced to work at an operating point in the nonlinear region of the characteristics with consequent distortion. Next, the AGC decreases the variable gain to the lower end of the available range. If the resultant degradation of the voltage resolution at the quantizer is not sufficient to compensate for the excessive signal strength, the ADC is saturated and no information is recoverable about the jamming waveform. Hereinafter, we overlook the quantization losses and limitations due to the AGC/ADC loop, which implies that our DoOs are necessarily devices equipped with sufficient bit resolution (e.g., between 8 and 14 bits) to take accurate TDOA measurements. In the following, the incoming signal is so ideally digitized into in-phase and quadrature (I/Q) raw samples, which are encoded with 64 bits each according to a *double-precision* floating point format. Instead, no assumption is made on the power level with respect to the *sensitivity* threshold of the receivers. Therefore, the interference might be received with a power spectral density buried under the front-end noise floor, hence with JNR lower than one. In our scenario, this situation may occur when the jammer is too weak or too far to be visible from all the DoOs.

  The implicit assumption of a sufficient bit resolution at the receivers is usually made in many papers dealing with jamming countermeasures (e.g., [116]), despite the major importance of the ADC quantization in reality. Since the methods proposed in this thesis are based on digitally processing the received signals, it is worth to clarify the requirements imposed to the specifications of DoOs. In fact, most of consumer-grade GNSS receivers feature 1/2-bit ADCs, which make them very likely to be saturated by the interference power received from jammers either too close or too powerful. But even when saturation does not occur, the incoming jamming waveforms are distorted by the digitization loop, thus frustrating any attempt of recovering the arrival times necessary to localize the jammers by trilateration. To avoid such an issue, the definition behind the DoOs mentioned throughout this chapter shall be restricted to exclude the entry-level receivers having low bit resolutions, which, however, are the majority of the devices on the consumer market. Examples of suitable and relatively inexpensive DoOs may be built with the SDR boards in Fig. 3.1, which are equipped with 8-bit ADCs. They could capture even powerful interference signals that are passed through an attenuator in the absence of the AGC. Likewise, other DoOs might mount robust GNSS receivers for safety critical applications, which have from 6 up to 14 quantization bits (e.g., that in [144, 145]). Therefore, when crowd-sourcing records of received interference samples from DoOs as proposed in [116], it is necessary a server-side selection of the snapshots that come from quantizers that are not saturated and have enough resolution. For an insight into the effect of poor resolutions (i.e., 2/3/4 bits) on the jammer localization performance with DRSS measurements, the reader may refer to [119].

- The GNSS signals are absent. – This assumption allows us to neglect the near-far problem that would arise when the JNR is negative and comparable to the power received from the satellites. A simple solution is presented in [128].

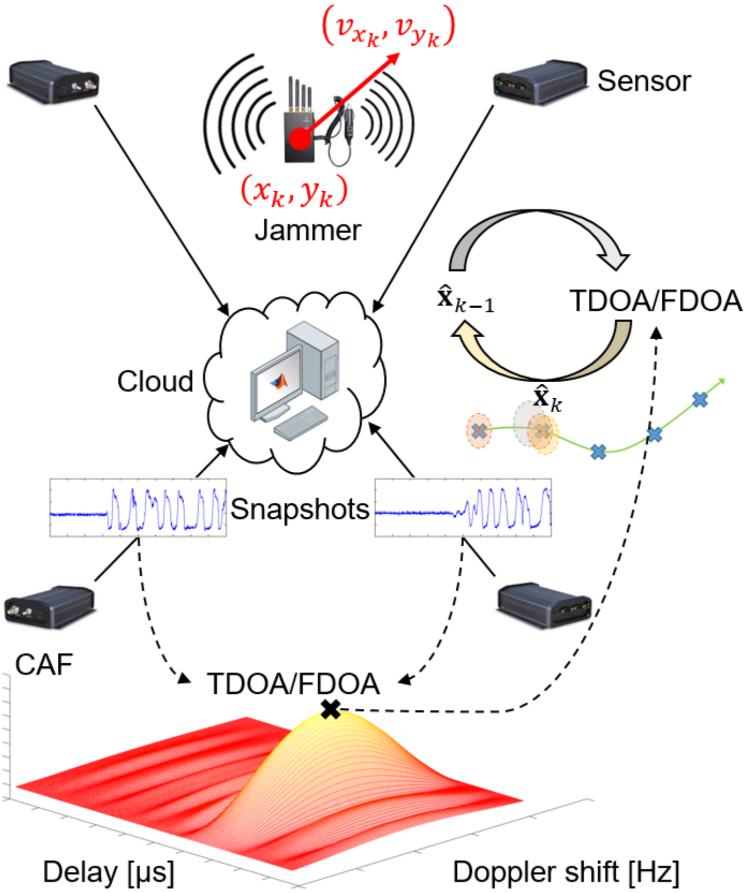

FIGURE 3.2: Application scenario where a jammer is tracked through simultaneous interference snapshots.

3.2.1 Maximum-Likelihood Estimation of TDOA/FDOA Pairs

Before being filtered and digitized and after being down-converted to baseband, the complex envelope of the analog jamming waveform $d(t)$ with respect to the time variable t might be formulated as a train of linear chirps:

$$d(t) = \sqrt{P_d(t)} \sum_{r=-\infty}^{\infty} \psi(t-(r-1)T_0) \tag{3.1}$$

where $P_d(t)$ is the interference power at the antenna input and T_0 is the repetition period in the sawtooth function $\psi(t)$, which is defined as

$$\psi(t) = \psi(t+T_0) = \begin{cases} \exp\left(j\pi(-B_{\mathrm{RF}}t + \frac{B_{\mathrm{RF}}}{T_0}t^2)\right), & 0 \leq t < T_0 \\ 0, & \text{otherwise} \end{cases} \tag{3.2}$$

with B_{RF} denoting the chirp passband (i.e., dual-sided) bandwidth centered around the carrier frequency f_c at RF, which usually coincide with that of the GNSS signal. This formulation is fairly accurate, because realistic nonlinearities due to the clock drift and the power amplifier do not actually compromise the linear rate of the chirps. One may verify this fact from the spectrograms experimentally measured at the output of the commercial jammers in [146]. Next, we may include into the model the presence of complex noise by adding a white Gaussian stochastic process

$v(t)$ with zero mean, as usual. Thence, the noisy signal at the input of the ADC is

$$s(t) = (d(t) + v(t)) * h_{\text{LB}}(t) \tag{3.3}$$

where $h_{\text{LB}}(t)$ is the impulse response of the front-end lowpass filter, which limits the visible spectrum of $d(t)$. Within every repetition period, the JNR is equal to

$$\rho = \frac{\frac{1}{T_0}\int_0^{T_0} |d(t) * h_{\text{LB}}(t)|^2 dt}{\frac{N_0}{2}\int_{-\infty}^{\infty} h_{\text{LB}}^2(t) dt} = \frac{2}{T_0 N_0 B_{\text{eq}}}\int_0^{T_0} P_d(t) dt = \frac{2\langle P_d(t)\rangle_{T_0}}{N_0 B_{\text{eq}}} \tag{3.4}$$

where B_{eq} is the equivalent noise passband bandwidth of the front-end filter, and N_0 is the single-sided noise power spectral density.
Owing to the need for a timely localization, we can determine the relative delay and Doppler frequency shift corresponding to a single jammer by examining the CAF. Let us first suppose that two noisy complex envelopes $s_0(t)$ and $s_1(t)$ are *simultaneously* received at two separate locations and only contain a common jamming waveform $d(t)$, which has traveled along the direct propagation path in the *line of sight*. We might then reformulate the jammed signals in the form of Eq. 3.3 to include distinct complex attenuation factors $A(t)$, delays $\tau(t)$ and frequency shifts $\nu(t)$, which depend on the distance and the velocity of the jammer with respect to the two fixed receivers, as follows

$$s_0(t) = A_0(t) d(t - \tau_0(t)) \exp(j2\pi\nu_0 t) + n_0(t) \tag{3.5}$$

$$s_1(t) = A_1(t) d(t - \tau_1(t)) \exp(j2\pi\nu_1 t) + n_1(t) \tag{3.6}$$

where

$$y_\tau(t) = \tau_1(t) - \tau_0(t) \tag{3.7}$$

$$y_\nu(t) = \nu_1(t) - \nu_0(t) \tag{3.8}$$

are the TDOA and the FDOA, respectively. The equations above entail the *narrowband approximation*: the Doppler shift is assumed as constant over the whole signal frequency range. This assumption is usually close to reality when the bandwidth is narrow compared to the carrier frequency, which is the case of interference on the GNSS bands. Therefore, if the complex envelopes in Eqs. 3.5 and 3.6 are down-converted by receivers perfectly synchronized in time and frequency, their continuous-time CAF is defined by

$$\text{CAF}(\Delta\tau, \Delta\nu) = \int_t^{t+T_{\text{obs}}} s_0(t') s_1^*(t' + \Delta\tau) \exp(j2\pi\Delta\nu t') dt' \tag{3.9}$$

where T_{obs} is the integration interval of the cross-correlation and here referred to as *observation* time. For legitimately invoking the narrowband approximation, the realistic delay spread caused by the jammer motion and the clock drift rate should have a negligible effect on the correlation functions over the observation. As previously anticipated, the CAF is the baseline for the joint estimation of the TDOAs and FDOAs produced by the jammer trajectory, which here play the role of the observables. Elaborating on the ML estimator in [147], we might extract the time averages of the TDOA and FDOA as the the peak arguments of the CAF magnitude:

$$\{\hat{y}_\tau, \hat{y}_\nu\} \approx \{\langle y_\tau(t)\rangle_{T_{\text{obs}}}, \langle y_\nu(t)\rangle_{T_{\text{obs}}}\} \approx \arg\max_{\Delta\tau,\Delta\nu} |\text{CAF}(\Delta\tau, \Delta\nu)| \tag{3.10}$$

where the approximation is due to the possible variations in time of the delays and Doppler frequency shifts, when the jammer is moving. When only one jammer is present, as we assumed, the CAF features a unique main lobe that corresponds to the region where most of the waveform energy is actually concentrated. If this jammer is also stationary, then the arguments of the peak are constant and so the equality in Eq. 3.10 becomes exact. Under these favorable circumstances, the ML estimator is efficient if the signal strength after cross-correlation greatly exceeds the additive Gaussian noise. In more detail, given identical front ends, the SNR at the output of the CAF is equal to

$$\gamma = \rho_{\text{eff}} T_{\text{obs}} B_{\text{eq}} \tag{3.11}$$

where $T_{\text{obs}} B_{\text{eq}}$ is the time-bandwidth product, also known as coherent processing gain, and

$$\frac{1}{\rho_{\text{eff}}} = \frac{1}{2}\left(\frac{1}{\rho_1} + \frac{1}{\rho_2} + \frac{1}{\rho_1 \rho_2}\right) \tag{3.12}$$

defines the effective JNR that combines the levels at the inputs of the two receivers. Therefore, the joint estimation in 3.10 is optimal if $\gamma \gg 1$ (e.g., higher than 10 dB). Since the receiver bandwidths is practically fixed, the integration time should be sufficiently long and the JNRs high enough, so that the observables tend to be mutually uncorrelated and independent of the observation location, even though they are related to the same jammer. When the scenario complies with these conditions, the unbiased ML estimator asymptotically achieves the CRB. As far as the the derivation of this bound is concerned, we may consider a couple of reasonable simplifications: the jammer energy spectral density is constant within T_{obs} ($P_d(t) = \langle P_d(t) \rangle_{T_{\text{obs}}}$) and flat over the receiver bandwidth B_{eq} ($B_{\text{RF}} \geq B_{\text{eq}}$). According to [147], the CRB is then defined by the following diagonal matrix

$$\mathbf{CRB} = \begin{bmatrix} \sigma_\tau^2 & 0 \\ 0 & \sigma_\nu^2 \end{bmatrix} \tag{3.13}$$

in which

$$\sigma_\tau = \frac{\sqrt{3}}{\pi B_{\text{eq}} \sqrt{\gamma}} \tag{3.14}$$

and

$$\sigma_\nu = \frac{\sqrt{3}}{\pi T_{\text{obs}} \sqrt{\gamma}}. \tag{3.15}$$

From Eq. 3.14 shows that TDOA measurements are especially precise when localizing a source of wideband signals, such as a jammer. Although commonly used in the research literature, the ideal diagonality of this definition of CRB might omit some strong correlations induced by specific non-stationary signals. This fact is plain to see with the example of a linear chirp: a delayed version of such a waveform looks the same as the one shifted by an equivalent negative Doppler frequency. Therefore, intuitively, off-diagonal correlation terms do exist among the estimates of TDOAs and FDOAs. A new derivation of the CRB for arbitrary deterministic signals is presented in [148] to account for these correlations. As for the sawtooth pattern that modulates the jamming waveforms in Eq. 3.1, the consequent correlation is instead neglected in the the classical CRB of Eqs. 3.13-3.15. Furthermore, this bound inherits another simplification from the model adopted for the measurement noise. The model in Eq. 3.3 is valid when measurements are perturbed by non-systematic errors due to thermal noise only. In reality, the accuracy of the observables depends also on the stability of all the stages in the front-end processing chain. For instance, contrary

to TDOAs, the FDOAs are measured with more precision for narrowband signals. Indeed, any modulation in frequency disrupts the periodicity of inner waveforms. These aspects are beyond the scope of our analysis.
The previous discussion considers analog signals. For the sake of simplicity, we can approximate the digital signal $s[n]$ as a sequence of complex values uniformly-sampled from $s(t)$ at time instants with indices n. By doing so, we do not take into account the quantization losses and the variable gain of the AGC/ADC loop. This approximation is legitimated by the negligible machine roundoff error of the 64-bit floating-point precision, which is among the assumptions made. This means that the raw I/Q samples can be expressed as the output of an ideal ADC by

$$s[l] \approx s(lT_s) \tag{3.16}$$

with sampling period T_s at the time instant numbered by the integer index l. As usual, the stream rate is supposed to fulfill the Nyquist criterion to avoid aliasing effects:

$$f_s = \frac{1}{T_s} \geq B_{\text{eq}}. \tag{3.17}$$

In order to have a discrete-time formulation suitable to signals digitized according to Eq. 3.16, we may discretize Eq. 3.9 as follows

$$\text{CAF}_l[n,p] = \sum_{m=l}^{l+N-1} s_0[m]s_1^*[m+n]\exp(j2\pi\frac{p}{N}m) \tag{3.18}$$

for fractional delays and normalized digital frequency shifts identified by the indices $n = 0, 1, ..., N-1$ and $p = -N/2, ..., N/2-1$, respectively, with $N = f_s T_{\text{obs}}$. This complex-valued function quantizes the delay-Doppler domain with a grid of equally-spaced points

$$Q = \left\{0, ..., T_{\text{obs}} - \frac{1}{f_s}\right\} \times \left\{-\frac{f_s}{2}, ..., \frac{f_s}{2} - \frac{f_s}{N}\right\} \tag{3.19}$$

the values of which are computed by cross-correlating the samples recorded by two synchronous and separate receivers. The TF resolution of the grid is tied to the support of the CAF, which is proportional to f_s and the T_{obs}. Unless the TDOA/FDOA pair happens to lie exactly on one of the points of Q, the energy associated to the actual peak leaks to the neighbor points underlying the CAF. This *leakage* of energy introduces a bias in the ML estimation that is more significant on the frequency axis than the time one if $T_{\text{obs}} \ll 1$ s. Infinitesimally-fine resolution with finite sample sizes is possible by means of the super-resolution radar proposed in [137], which enables the distinction of multiple and closely-spaced signal sources in the delay-Doppler domain. Nevertheless, this technique employs a polynomial time in the number of targets and is conceived for noiseless (i.e., with very high SNR, or equivalently JNR) and clutter-free conditions. Similar approaches in [79, 136] further depend on the flatness of the incoming signal spectra. For the reasons previously explained, the CAF represents a much more robust and preferable choice for the application under study. More specifically, since our goal is a low-complexity and real-time estimation of TDOA/FDOA pairs, we can resort to the simplest and fastest implementation of the discrete CAF:

$$\text{CAF}_l[n,p] = \text{FFT}_{m\to p}\{s_0[m]s_1^*[m+n]\} \tag{3.20}$$

which is straightforward thanks to the circular convolution theorem for the DFT. Further computational savings are obtained by restricting the indices n and p over a rectangular subset of points

$$Q' = \{0, ..., y_{\tau_{\max}}\} \times \{-y_{\nu_{\max}}, ..., y_{\nu_{\max}}\} \tag{3.21}$$

where $y_{\tau_{\max}}$ and $y_{\nu_{\max}}$ are set according to geometric considerations about the receiver coverage and the maximum possible jammer speed. Likewise the ML estimator in Eq. 3.10, identifying the indices corresponding to the peak enables a rough estimate of the TDOA/FDOA pair. Beyond the limited TF resolution, a finer-grained estimation is achieved by interpolating the magnitude of the CAF among the nine points of the grid around the peak. Despite the interpolation, though, the aforementioned energy leakage due to quantization prevents the ML estimator from being optimal in Q. Near optimality is still achievable with one moving jammer, as long as none of the time-varying observables $y_\tau(t)$ and $y_\nu(t)$ change by more than one point in the grid during the integration time. We may formulate this constraint through the following inequalities:

$$|y_\tau(t + T_{\text{obs}}) - y_\tau(t)| < 1/fs \tag{3.22}$$

$$|y_\nu(t + T_{\text{obs}}) - y_\nu(t)| < fs/N \tag{3.23}$$

which impose the observation to be short enough compared to the swiftest variation possible of the TDOA/FDOA pair caused by the speedy motion of the jammer. Otherwise, the energy of the correlation in the delay-Doppler domain is smeared as a side effect of the dynamics, making any search for the peak prone to biases. While this constraint imposes a higher bound on T_{obs}, a lower bound is represented by the CRB. In fact, the integration time should be long enough to comply with the precision desired for the observables. In this regard, despite the presence of biases due to the effects of energy leakage and jammer motion in Q, we may still refer to the CRB for the "overoptimistic" performance of the theoretically unbiased ML estimator. The calculation of this bound from Eqs. 3.13-3.15 actually requires an a-priori knowledge of the JNR in Eq. 3.4, which otherwise shall be estimated. This estimation is possible through a simple signal-to-noise variance estimator, once the noise power is known and a jammer is detected. Often when initialized, the receiver is supposed to be free from interference to perform a calibration phase, in which the noise level of the environment could be estimated from the incoming stream of I/Q samples, and possibly used to later detect jammers. As far as the maximum of the CAF, other than interpolating the values of the FFT in Eq. 3.20, more advanced algorithms are presented in [147] for the two-step and computationally-efficient computation of the CAF. They are useful to process long sequences of samples with greater TF resolution and to compensate for the effects of dynamic observables.

Generally, the ML estimator described copes well with the presence of one jammer only, because the CAF features a unique main lobe in the quantized delay-Doppler domain. The widths of this correlation lobe are about $1/B_{\text{RF}}$ in time and $1/T_{\text{obs}}$ in frequency, if the jammer is continuously transmitting during the observation. Particularly, the spectral profile has a sharp peak of width $1/B_{\text{eq}}$, when the wideband sawtooth of the jamming waveform spans over the whole frequency range of the front-end filter (i.e. $B_{\text{RF}} \geq B_{\text{eq}}$). Otherwise, any bandwidth narrower than this range or any pulsed modulation in time smooth the main lobe, thus degrading the precision of the ML estimator. This undesirable effect is also predicted by the CRB. We may model this behavior by approximating the actual shape of the normalized CAF with a two-dimensional *sinc-like* surface:

$$\overline{\mathrm{CAF}}(\Delta\tau, \Delta\nu) \approx \mathrm{sinc}((\Delta\nu - y_\nu)T_{\mathrm{obs}})\,\mathrm{sinc}((\Delta\tau - y_\tau)B_{\mathrm{eq}})\mathrm{e}^{-j\pi(\Delta\nu - y_\nu)T_{\mathrm{obs}}} \tag{3.24}$$

for $B_{\mathrm{RF}} \geq B_{\mathrm{eq}}$, as argued in [140] and apparent in Fig. 3.3. As long the input JNRs and the processing gain are sufficient to discriminate the peak of energy associated to the TDOA/FDOA pair from the side lobes and the noise floor, the observables are *always* well identified with the arguments of the maximum magnitude of the CAF, regardless of the jamming waveform. This advantageous property is crucial, because *no a-priori information* about the signal structure is necessary to perform a reliable estimation. On the contrary, generally speaking, the same does not hold when two or more co-channel sources are simultaneously present or, equivalently, when replicas of the signal are received from the surrounding scatterers. With the reception of multiple jamming waveforms, indeed, the respective CAFs mutually interact in both *constructive* and *destructive* manners, depending on the relative distances and velocities of the jammers with respect to the receiver. Hence, the summation of their CAFs in the delay-Doppler domain might cause missed detections of the interference energy, due to the cancellation or superposition of main lobes, as well as false alarms, if significant peaks are produced by spurious energy. Such a mutual interaction is especially harmful when the jammers are visible with very different JNR levels and/or located close to each other. Consequently, the ML estimator is likely to be mislead and so potentially biased. As a matter of example, the issue of relying on the correlation peaks for localizing two closely-spaced jammers is exemplified in the CAFs depicted in Figs. 3.4a and 3.4b. This challenge is tackled in [140], where the TDOAs and FDOAs of multiple jammers are initially obtained as the solutions of an optimization problem, which adopts the model in Eq. 3.24. A successive clustering process combines these estimates over a grid that quantizes the possible locations in the area in order to find data associations and to mitigate false alarms and missed detections. While this algorithm works in the position domain, a completely different strategy could shift the problem of estimating and discriminating the TDOA/FDOA pairs from the delay-Doppler domain to the TF one, the analysis of which is introduced in Chapter 2. This innovative solution is demonstrated later in this chapter to be able of bypassing at once both the uncertainties affecting the ML estimator and the data association.
An interesting aside from the current topic regards the CAF. The continuous-time formulation in Eq. 3.9 is a generalization of the ambiguity function (AF) originally defined by

$$\mathrm{AF}(\Delta\tau, \Delta\nu) = \int_t^{t+T_{\mathrm{obs}}} s\left(t' + \frac{\Delta\tau}{2}\right) s^*\left(t' - \frac{\Delta\tau}{2}\right) \exp(j2\pi\Delta\nu t')dt' \tag{3.25}$$

where the $s(t)$ denotes either the complex envelope of the lowpass signal or the analytic form of the passband one. The AF and the Wigner-Ville distribution (WVD) mentioned in Chapter 2 are related to each other through the instantaneous ACF of the signal: the former one is defined as the FT with respect to delay variable, while the latter one is the FT with respect to time variable. Therefore, the AF evaluates the correlation over the delay-Doppler domain, whereas the WVD is a TF representation of the signal energy. Both of them suffer from the presence of artifacts, which result from their quadratic nature.

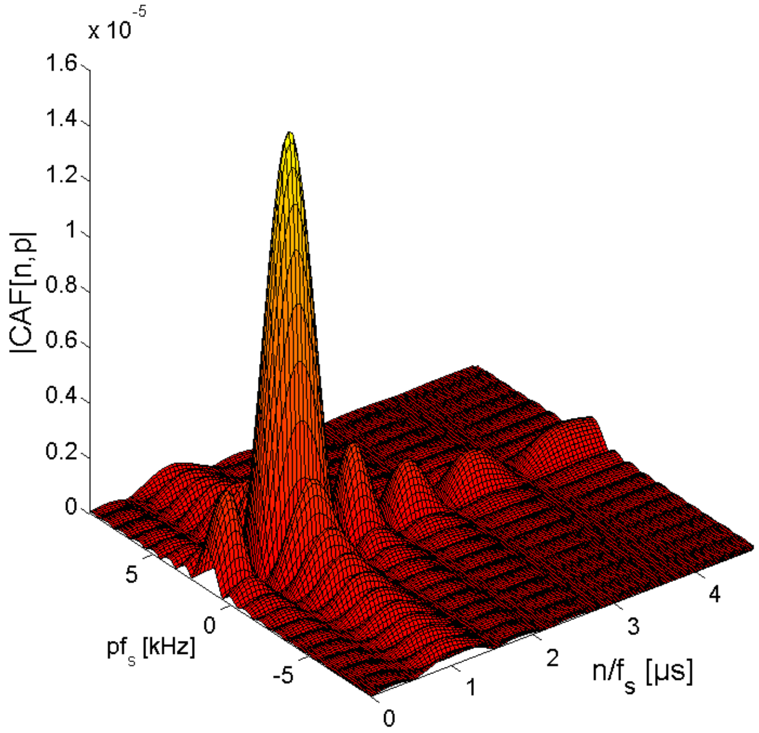

FIGURE 3.3: Sinc-like shape of the CAF computed from Eq. 3.20 for $f_s = 20$ MHz, $B_{\text{eq}} = 1$ MHz and $T_{\text{obs}} = 1$ ms in the presence of one jammer with $B_{\text{RF}} = 2$ MHz.

3.2.2 Nonlinear Estimation of the Jammer State

In the application scenario summarized by Fig. 3.2, we can count on a number M of DoOs that concurrently record sequences of I/Q raw samples. These recordings are snapshots containing a common jamming waveform $d(t)$ in the form of Eqs. 3.5 and 3.6. They are collected over a time span of one observation (i.e., T_{obs}) and sent to the cloud. Here, a central processor computes $M-1$ discrete CAFs by separately cross-correlating the samples $s_0[n]$ received by one DoO chosen as *reference* with each of the snapshots $s_i[n]$ captured by the other $M-1$ devices, which are indexed with $i = 1, ..., M-1$. Given any of the couples of snapshots $\{(0, i) \mid i = 1, ..., M-1\}$, the FFT in Eq. 3.20 is performed for every frequency index p in the subset of points Q' that quantize the ambiguity region of interest in the delay-Doppler domain. From the resultant CAF, the TDOA/FDOA pair $\{y_{\tau_i}, y_{\nu_i}\}$ of the jamming waveform is identified as the arguments of the maximum magnitude found by interpolation, according to the ML estimator in Eq. 3.10. These estimates are independent of the choice of the reference DoO, as proven in [55]. Overall, the estimation returns a column vector of $2(M-1)$ observables:

$$\mathbf{y} = [\mathbf{y}_\tau \, \mathbf{y}_\nu]^T \tag{3.26}$$

where $\mathbf{y}_\tau = [y_{\tau_1} \, ... \, y_{\tau_{M-1}}]$ and $\mathbf{y}_\nu = [y_{\nu_1} \, ... \, y_{\nu_{M-1}}]$ are retrieved for every time epoch in which a dataset of snapshots is made available to the cloud by the DoOs. A time epoch coincides with interval of duration T_{obs}. The sequence of epochs is numbered by an integer index k, such that each of these intervals consist of N samples with

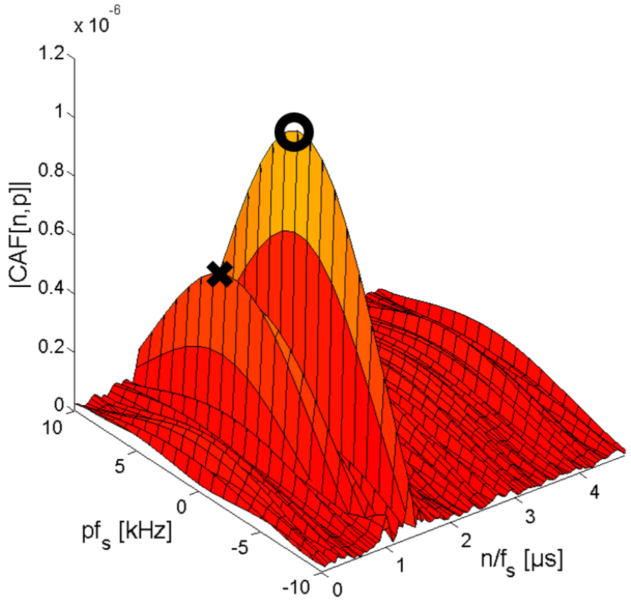

(A) The two jammers are separated by about 110 m.

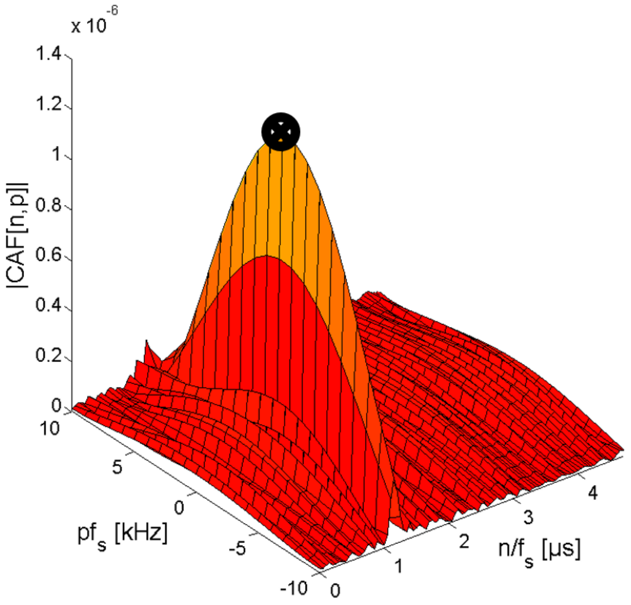

(B) The two jammers are separated by about 28 m.

FIGURE 3.4: Examples of CAF magnitudes computed from Eq. 3.20 for $f_s = 20$ MHz, $B_{\text{eq}} = 5$ MHz and $T_{\text{obs}} = 0.1$ ms in the presence of two jammers with $B_{\text{RF}} = \{10 \text{ MHz}, 20\text{MHz}\}$.

indices $l = (k-1)N+1, ..., kN$. For a matter of power and bandwidth consumptions, the devices used as receivers cannot transfer data continuously to the central processing unit, but in actuality they rather transmit intermittently. Without loss of generality, let us assume that this intermittence is periodic with snapshots aligned and equally-distributed in time. For the reasons mentioned above, the period between consecutive datasets is supposed to be longer than just a single observation and, more specifically, defined by $T_{obs}R_s^{-1}$, where $R_s \leq 1$ is the snapshot *rate* normalized on the integration time. The break interval between successive transmissions,

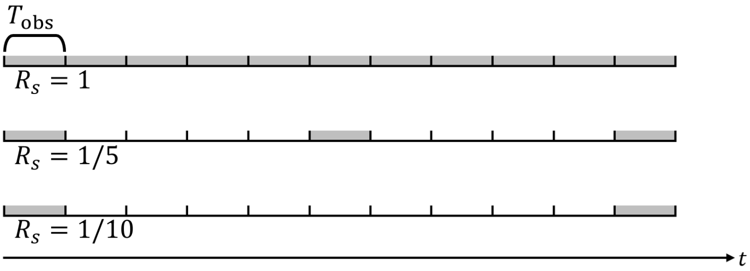

FIGURE 3.5: Normalized snapshot rates with periods scaled by integer factors.

over which the DoOs stay idle, is as long as $T_{obs}(R_s^{-1} - 1)$. A unitary rate then corresponds to a continuous tracking, seamlessly epoch after epoch, without breaks. Hereinafter, for the sake of simplicity, we vary the density of snapshots over time by scaling the period R_s^{-1} by integer factors, as shown in Fig. 3.5. By doing so, we can number the sequence of datasets through the same index k used to count the epochs. Consequently, this enumeration applies as well to the vectors $\mathbf{y}_k$ in the form of Eq. 3.26 that collect the TDOA/FDOA pairs estimated from the snapshots. In detail, given a fixed R_s, the sequence of observables may be expressed as $\left(\mathbf{y}_1, ..., \mathbf{y}_k, ..., \mathbf{y}_K \mid \bmod\left(k-1, R_s^{-1}\right) = 0\right)$ for a number of epochs $K \geq 1$, from which follows that the entire sample stream is made of $L = KN$ samples with $l = 1, ..., L$. The pace of these observations should be high enough to track a jammer rapidly maneuvering. Thereupon, in the following, the origin of the jamming attempt is supposed to move without accelerating along a straight trajectory, which does not rotate. This simplification is not necessarily surreal, because the track of any maneuver is theoretically decomposable into constant-speed segments, if the observables are dense enough in time.

Now, let us consider the unknown position and velocity of the jammer at the time epoch k in $\mathbf{x}_k$ as the hidden state of a stochastic process possessing Markov property onto an infinite and uncountable space. This state can be recursively inferred by combining the previous estimates $\hat{\mathbf{x}}_{k-1}$ and the current observables $\mathbf{y}_k$. Within this recursion, the TDOA/FDOA pairs represents measurements that are indirectly related to $\mathbf{x}_k$ by

$$\mathbf{y}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{v}_k \tag{3.27}$$

where $\mathbf{v}_k$ is the additive measurement noise. If the equipment is properly calibrated and does not induce further errors, then $\mathbf{v}_k$ is a multivariate Gaussian random variable with zero mean and known and positive-definite covariance matrix $\mathbf{R}_k$, which shall be coherent with considerations made about the hypothetical accuracy of the ML estimator. The formulation of $\mathbf{h}$ depends on whether the observable is in time or frequency. For reasons of clarity, we omit the subscript k in the following equations. The function h_{τ_i} relates the *exact* (i.e., noiseless) TDOA scalar y_{τ_i} to the positions of the jammer and two DoOs according to

$$\begin{aligned} y_{\tau_i} &= h_{\tau_i}(\mathbf{x}) = (d_0 - d_i)/\mathrm{c} = \\ &\frac{\left(\sqrt{(x-x_0)^2 + (y-y_0)^2} - \sqrt{(x-x_i)^2 + (y-y_i)^2}\right)}{\mathrm{c}} \end{aligned} \tag{3.28}$$

where c is the speed of light, (x_0, y_0) is the location of the device chosen as reference, (x_i, y_i) is the location of any other receiver, and d_0 and d_n are the respective distances from the jammer of these two. Similarly, the function h_{ν_i} relating the exact FDOA scalar y_{ν_i} to $\mathbf{x}$ is

$$\begin{aligned} y_{\nu_i} &= h_{\nu_i}(\mathbf{x}) = f_c(v_{r_0} - v_{r_i})/\mathrm{c} = \\ &\frac{f_c}{\mathrm{c}}\left(\frac{v_x(x-x_0)+v_y(y-y_0)}{d_0} - \frac{v_x(x-x_i)+v_y(y-y_i)}{d_i}\right) \end{aligned} \tag{3.29}$$

where v_{r_0} and v_{r_i} are the radial velocity components of the jammer with respect to the reference DoO and any of the others. To ensure the numerical stability, we combine Eqs. 3.28 and 3.29 into a nonlinear vector function defined by

$$\mathbf{h} = [\mathrm{c}\,\mathbf{h}_\tau\,\mathbf{h}_\nu]^T \tag{3.30}$$

with $\mathbf{h}_\tau = [h_{\tau_1} \ldots h_{\tau_{M-1}}]$ and $\mathbf{h}_\nu = [h_{\nu_1} \ldots h_{\nu_{M-1}}]$. The measurement error covariance becomes then a well-conditioned diagonal matrix of order $2(M-1)$

$$\mathbf{R} = \mathrm{diag}\,(\mathrm{c}\,\sigma^2_{\tau_1}, \ldots, \mathrm{c}\,\sigma^2_{\tau_{M-1}}, \sigma^2_{\nu_1}, \ldots, \sigma^2_{\nu_{M-1}}) \tag{3.31}$$

the elements of which are the precisions at the CRB defined by Eqs. 3.14 and 3.15 for every i-th couple of receivers as a function with respect to the time-varying JNR estimate. Accordingly, we redefine Eq. 3.26 as

$$\mathbf{y} = [\mathrm{c}\,\mathbf{y}_\tau\,\mathbf{y}_\nu]^T. \tag{3.32}$$

Once the index k is restored into the notation, we might notice that the matrix $\mathbf{R}_k$ changes from epoch to epoch, as the output SNR levels denoted by γ_{i_k} varies over time with the evolving location of the jammer.

Retrieving $\mathbf{x}_k$ from $\mathbf{y}_k$ is a nonlinear estimation problem, which shall be preferably solved before the next observables (i.e., at $k + R_s^{-1}$) are made available. In order to keep up with the jammer motion in real time, fast resolution methods of the TDOA and FDOA equations generally resort to a Taylor-series expansion of $\mathbf{h}$. The first-order approximation retains the first two terms of this series and is so the most computationally efficient, but the least accurate. In other words, it *linearizes* the highly nonlinear measurements curves (i.e., hyperbolic for the TDOAs) about the an a-priori rough estimate $\hat{\mathbf{x}}_{0_k}$, as follows

$$\mathbf{y}_k \approx \mathbf{h}(\hat{\mathbf{x}}_{0_k}) + \mathbf{H}_k(\mathbf{x}_k - \hat{\mathbf{x}}_{0_k}) + \mathbf{v}_k \tag{3.33}$$

where $\mathbf{H}_k$ is the Jacobian matrix of $\mathbf{h}$ at the current time epoch:

$$\mathbf{H}_k = \begin{bmatrix} \mathrm{c}\,\nabla\mathbf{h}_\tau(\mathbf{x}_k) \\ \nabla\mathbf{h}_\nu(\mathbf{x}_k) \end{bmatrix} = \begin{bmatrix} \mathrm{c}\,\frac{\partial h_{\tau_1}(\mathbf{x}_k)}{\partial x} & \mathrm{c}\,\frac{\partial h_{\tau_1}(\mathbf{x}_k)}{\partial y} & \mathrm{c}\,\frac{\partial h_{\tau_1}(\mathbf{x}_k)}{\partial v_x} & \mathrm{c}\,\frac{\partial h_{\tau_1}(\mathbf{x}_k)}{\partial v_y} \\ \vdots & \vdots & \vdots & \vdots \\ \mathrm{c}\,\frac{\partial h_{\tau_1}(\mathbf{x}_k)}{\partial x} & \mathrm{c}\,\frac{\partial h_{\tau_{M-1}}(\mathbf{x}_k)}{\partial y} & \mathrm{c}\,\frac{\partial h_{\tau_{M-1}}(\mathbf{x}_k)}{\partial v_x} & \mathrm{c}\,\frac{\partial h_{\tau_{M-1}}(\mathbf{x}_k)}{\partial v_y} \\ \frac{\partial h_{\nu_1}(\mathbf{x}_k)}{\partial x} & \frac{\partial h_{\nu_1}(\mathbf{x}_k)}{\partial y} & \frac{\partial h_{\nu_1}(\mathbf{x}_k)}{\partial v_x} & \frac{\partial h_{\nu_1}(\mathbf{x}_k)}{\partial v_y} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{\partial h_{\nu_{M-1}}(\mathbf{x}_k)}{\partial x} & \frac{\partial h_{\nu_{M-1}}(\mathbf{x}_k)}{\partial y} & \frac{\partial h_{\nu_{M-1}}(\mathbf{x}_k)}{\partial v_x} & \frac{\partial h_{\nu_{M-1}}(\mathbf{x}_k)}{\partial v_y} \end{bmatrix} \tag{3.34}$$

where the elements in the $2(M-1)$ rows and 4 columns are expressed as

$$\frac{\partial h_{\tau_i}(\mathbf{x}_k)}{\partial x} = \frac{x_k - x_0}{\mathrm{c}\, d_{0_k}} - \frac{x_k - x_i}{\mathrm{c}\, d_{i_k}} \tag{3.35}$$

$$\frac{\partial h_{\tau_i}(\mathbf{x}_k)}{\partial v_x} = \frac{\partial h_{\tau_i}(\mathbf{x}_k)}{\partial v_y} = 0 \tag{3.36}$$

$$\frac{\partial h_{\nu_i}(\mathbf{x}_k)}{\partial x} = \frac{f_c}{\mathrm{c}}\Big(\frac{v_{x_k} d_{0_k}^2 - r_{v_{0_k}} d_{0_k}(x_k - x_0)}{d_{0_k}^3} - \frac{v_{x_k} d_{i_k}^2 - r_{v_{i_k}} d_{i_k}(x_k - x_i)}{d_{i_k}^3}\Big) \tag{3.37}$$

$$\frac{\partial h_{\nu_i}(\mathbf{x}_k)}{\partial x} = \frac{f_c}{\mathrm{c}}\Big(\frac{x_k - x_0}{d_0} - \frac{x_k - x_i}{d_i}\Big). \tag{3.38}$$

Any error minimization based on this approximation is sub-optimal and does not guarantee the convergence to the global minimum. Indeed, under some circumstances where the transmitter-receivers geometry is poor, the significant errors with respect to the true values of the nonlinear function can lead the estimator to a local minimum or make it diverge. The proximity of the initial estimate $\hat{\mathbf{x}}_{0_k}$ to $\mathbf{x}_k$ is important to ensure the convergent behavior of the estimator. A good estimate may be guessed from prior information about the scenario or, in turn, initialized with a sub-optimal procedure. By looking at the partial derivatives of $\mathbf{H}_k$, some consideration can be made. Since the range terms at the numerator and denominator in Eq. 3.35 are of the same order of magnitude, the TDOAs are less sensitive to the distance between transmitter and receiver than AOA and DRSS measurements. Therefore, as far as the observation of differential delays is concerned, the DoOs could have sparse locations spread over large areas, without adverse consequences. A critical situation occurs when the TDOA is observed by a device too close to the jammer (i.e., $d_{i_k} \approx 0$), as the precision tends to drop. Both these aspects are evident also for the FDOAs in Eq. 3.38. The effect of differential Doppler frequency shifts on the position estimation is evaluated by 3.37. For this impact to be meaningful, the relative velocity at play should be large enough to outweigh the quantization losses in the discrete CAF *after* interpolation. For instance, if TF resolution of the delay-Doppler domain is insufficient to measure the average speed of the jammer, the addition of FDOA is likely to impair the overall localization performance, rather than enhancing it. In this case, the employment of solely TDOA measurements is more advisable. Besides, the errors realistically induced by the equipment could further worsen this detrimental impact. Last but not least, Eq. 3.36 means that the present formulation neglects the delay-Doppler *coupling* effect typical for chirp-like waveforms in the presence of high relative velocities. Simply put, the top speed of the jammer does not affect the TDOA observables.
The simplest approach based on Eq. 3.33 makes use the LS estimator in [50]:

$$\hat{\mathbf{x}}_k = \begin{cases} \hat{\mathbf{x}}_{0_k} + (\mathbf{H}_k^T \mathbf{R}_k^{-1} \mathbf{H}_k)^{-1} \mathbf{H}_k^T \mathbf{R}_k^{-1} (\mathbf{y}_k - \mathbf{h}(\hat{\mathbf{x}}_{0_k})), & \mod (k-1, R_s^{-1}) = 0 \\ \hat{\mathbf{x}}_{0_k} & \text{otherwise} \end{cases} \tag{3.39}$$

that can be made recursive by linearizing about the previous estimate with

$$\hat{\mathbf{x}}_{0_k} = \hat{\mathbf{x}}_{k-1}. \tag{3.40}$$

This estimation method may be regarded as optimal from the perspective of the linearized $\mathbf{h}$, but it does inherit a bias from the first-order approximation of the actual TDOA/FDOA curves. The downside is that a new $\hat{\mathbf{x}}_k$ is processed only when up-to-date datasets are transfered by the DoOs. While these devices are idle, no adjustment is made on the state estimate, which then becomes outdated in the absence of measurements. As a consequence, the error tends to increase proportionally to the time elapsed since the last snapshots and the jammer speed. To keep tracking $\hat{\mathbf{x}}_k$ despite the momentary lack of observables, we may resort to a model of the differential equation of the target motion to predict the evolution of the state. This model is used either as aid or backup, depending whether $\mathbf{y}_k$ is available or not. This alternative approach is enabled by many Bayesian methods, such as the well-known KF. This information filter propagates the knowledge about $\mathbf{x}_k$ through a recursion consisting of two steps: state *prediction* and measurement *update*. At every time epoch, the recursive *posterior* estimate $\hat{\mathbf{x}}_k$ is a statistically weighted average between the *prior* state $\hat{\mathbf{x}}_k^-$ predicted and, whenever possible, the innovation residual (e.g., $\mathbf{y}_k - \mathbf{H}_k\hat{\mathbf{x}}_k^-$). The weighted is the so-called Kalman gain. As for the former step, the uncertainty on $\mathbf{x}_k$ undergoes a dynamics model that we may formulate as follows

$$\mathbf{x}_k^- = \mathbf{A}\mathbf{x}_{k-1} + \mathbf{w}_k \tag{3.41}$$

where $\mathbf{A}$ is the time-invariant transition matrix associated to a stationary and constant-velocity trajectory:

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & T_{\text{obs}} & 0 \\ 0 & 1 & 0 & T_{\text{obs}} \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{3.42}$$

which should well approximate any maneuver of the jammer throughout a period as long as $T_{obs}R_s^{-1}$. The column vector $\mathbf{w}_k$ is additive state noise, which we represent as a multivariate Gaussian random variable with zero mean and covariance matrix $\mathbf{Q}$. This noise essentially scales the extent of a fading memory effect. Suitable elements of $\mathbf{Q}$ shall be set to mirror the degree of confidence into the model. This matrix should always have a non-zero determinant to compensate for roundoff errors, even though Eq. 3.41 might match the reality. On the contrary, if the model is unreliable, a high state noise should lessen the influence of $\hat{\mathbf{x}}_k^-$ on $\hat{\mathbf{x}}_k$ in order to force the KF to rely solely on $\mathbf{y}_k$, similarly to the LS estimator. Although modeling the dynamics generally provides an advantage over the LS, the choice of wrong values for the elements of $\mathbf{Q}$ could end up harming the stability of the KF, thus resulting in a degradation of the tracking performance, instead of an enhancement. In practice, we could adapt this matrix epoch-by-epoch to the time-varying uncertainty on the model as well as the time pasted without snapshots to process, for the achievement of quasi-optimal performance. For the application being studied, the description of the dynamics in Eq. 3.41 is trustworthy and used to predict $\hat{\mathbf{x}}_k^-$ from the previous a-posteriori estimate $\hat{\mathbf{x}}_{k-1}$. In the absence of new pairs of TDOA/FDOA measurements, the current posterior estimate is simply obtained by forwarding the prior one as

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}_k^-, \qquad \operatorname{mod}\left(k-1, R_s^{-1}\right) \neq 0 \tag{3.43}$$

that is output of the linear model in Eq. 3.41. The same happens for the estimation of the error covariance:

$$\hat{\mathbf{P}}_k = \hat{\mathbf{P}}_k^- = \mathbf{A}\hat{\mathbf{P}}_{k-1}\mathbf{A}^T + \mathbf{Q}, \qquad \operatorname{mod}\left(k-1, R_s^{-1}\right) \neq 0 \tag{3.44}$$

where $\hat{\mathbf{P}}_k$ and $\hat{\mathbf{P}}_k^-$ denote the posterior and the prior matrices, respectively. Otherwise, whenever datasets are finally made available to the processor, $\hat{\mathbf{x}}_k^-$ is employed as input into a second step that updates $\hat{\mathbf{x}}_k$, according to the chosen approximation of the nonlinear model in Eq. 3.27. In this regard, there exist several variants of the KF, which feature different equations to deal with nonlinearities in either the dynamics or the observations or both. Among the conventional ones, we focus on the first-order extended KF (EKF1), the second-order EKF (EKF2), and the unscented KF (UKF). At every epoch, the EKF1 first approximates the PDF of the state with the respective prior mean and covariance and propagates them through the linear dynamic model. Secondly, only at those epochs where observables are present, this knowledge goes through the linearized observation in Eq. 3.33. At end of these two steps, the resultant a-posteriori estimate of the mean and covariance of the state distribution are not guaranteed to be close to the actual ones. Nonetheless, if the linearization, the dynamic model, and the Gaussian measurement and state noises provide a truthful representation of the actuality, this filter converges to the estimate that minimizes the MSE. Conversely, if a significant bias exists due to the poor approximation of the TDOA and FDOA curves, the EKF2 extends the EKF1 by retaining one more term of the Taylor-series expansion of $\mathbf{h}$. This modification potentially makes this second filter more unstable and requires additional complexity for the calculation of the Hessian matrix $\mathbf{G}_k$, which contains the second-order partial derivatives over $2(M-1)$ rows and 10 columns. The derivation of $\mathbf{G}_k$ from $\mathbf{H}_k$ is straightforward and here omitted for the sake of brevity. The UKF belongs to the family of sigma-point KFs, which is built around the idea that it is easier to approximate a probability distribution than a nonlinear function. In fact, it improves on the EKF1 and EKF2 by propagating the PDF of the state through a minimal set of deterministically-chosen and so-called sigma points, rather than just the mean and the covariance estimates. These points are weighted and scaled, such that they match the first- and second-order moments of the prior state distribution. They are generated through the scaled unscented transform, which allows for accurately tracking the statistics of a random variable that undergoes a nonlinear transformation, such as that of $\mathbf{h}$. Once passed through the nonlinearity, the sigma points capture the posterior mean and covariance of the state to at least the second order of accuracy (i.e., the same of the EKF2) for any nonlinearity. Another desirable property of the UKF lies in the computational expense, which is on the same order of the EKF1 and also lacks of any explicit calculation of the gradients of $\mathbf{h}$ (i.e., $\mathbf{H}_k$ and $\mathbf{G}_k$).

An historical perspective from the LS method to the original KF is given [65]. The scientific literature offers tens of papers reviewing or applying the KF in diverse fields. Among them, the motivations and the developments behind the UKF are illustrated in [149]. A variety of sigma-point schemes can be combined with the unscented transform to address the deficiencies of the linearization. For instance, in [150], this transform is also used to derive a derivative-free version of the EKF2. The relations between the EKF1/EKF2 and the UKF are extensively investigated in [151]. Depending on the nonlinear functions that underlie the estimation problem, these two latter papers demonstrate how the EKF2 could outperform the UKF, which is otherwise the standard solution. Aside from the nonlinearity, the limits of a single KF versus a more versatile IMM estimator are quantified as a function of the target maneuverability in [152]. The implementation of multiple models is certainly a necessary and future development for the application of interest in order to relax the demand for high snapshot rate when the jammer can accelerate and turn.

The amount of knowledge accumulated or, equivalently, uncertainty left about the $\mathbf{x}_k$ is fully characterized by the Fisher information matrix (FIM) that is evaluated over the joint unconditional PDF $p([\mathbf{x}_0 \dots \mathbf{x}_k], [\mathbf{y}_1 \dots \mathbf{y}_k])$ for the whole collections of state and measurement vectors up to the time epoch k. According to the recursive formulation of this matrix in [153], we may describe the information extrapolated by the EKF1 with

$$\mathbf{J}_k = \mathbf{Q}^{-1} + \mathbf{H}_k^T \mathbf{R}_k^{-1} \mathbf{H}_k - (-\mathbf{A}^T \mathbf{Q}^{-1})^T (\mathbf{J}_{k-1} - \mathbf{A}^T \mathbf{Q}^{-1} \mathbf{A})(-\mathbf{A}^T \mathbf{Q}^{-1}). \tag{3.45}$$

This equation provides an amount of information larger than that of the dynamicsless LS estimator and smaller than that granted by higher-order approximations, namely the EKF2 and the UKF. As such, it turns out to be useful for benchmarking the variants of the KF. In wider terms, the FIM is essential for assessing the sub-optimal performance of the nonlinear estimation by comparison with the theoretically best-case accuracy of the jammer state estimated. In this regard, one may refer to the inequality of the joint unconditional posterior (or Bayesian) CRB (PCRB) for biased estimators:

$$\mathbf{P}_k = \mathrm{E}_{p\left(\mathbf{x}_k, [\mathbf{y}_1 \cdots \mathbf{y}_k]\right)} \left[(\mathbf{x}_k - \hat{\mathbf{x}}_k)(\mathbf{x}_k - \hat{\mathbf{x}}_k)^T \right] \geq \mathbf{PCRB}_k = \mathbf{J}_k^{-1} \tag{3.46}$$

where $\mathbf{P}_k$ denotes the error covariance up to the k-th epoch, which is not to be confused with the k-th a-posteriori covariance matrix $\hat{\mathbf{P}}_k$ estimated by any KF. Other versions of the PCRB are reviewed in [154]. The recursion in Eq. 3.45 may be initialized with the initial error covariance matrix:

$$\mathbf{J}_0 = \hat{\mathbf{P}}_0^{-1} \tag{3.47}$$

which we set consistently with the geometry of the DoOs deployment and the area monitored. If the FIM is known, one may convert Eq. 3.46 into a lower bound on the estimation RMSE as follows

$$\mathrm{RMSE}_k = \sqrt{\mathrm{trace}(\mathbf{P}_k)} \geq \sqrt{\mathrm{trace}(\mathbf{PCRB}_k)}. \tag{3.48}$$

This inequality may be decomposed to discriminate between position and velocity components of the jammer state:

$$\mathrm{RMSE}_{k_{(x,y)}} \geq \sqrt{\mathrm{PCRB}_{11_k} + \mathrm{PCRB}_{22_k}} \tag{3.49}$$

$$\mathrm{RMSE}_{k_{(vx,vy)}} \geq \sqrt{\mathrm{PCRB}_{33_k} + \mathrm{PCRB}_{44_k}}.. \tag{3.50}$$

At each time epoch, the error (i.e., $\mathbf{x}_k - \hat{\mathbf{x}}_k$) is not actually a constant vector, because $\mathbf{x}_k$ changes while being observed as the jammer is moving within T_{obs}. Therefore, we shall redefine the error sample-by-sample as

$$\mathbf{e}_l = \mathbf{x}_k[l] - \hat{\mathbf{x}}_k, \quad l \in \{(k-1)N + 1, ..., Nk\} \tag{3.51}$$

so that we have a matrix

$$\mathbf{x}_k - \hat{\mathbf{x}}_k = [\mathbf{e}_{(k-1)N+1} \cdots \mathbf{e}_{kN}]. \tag{3.52}$$

We calculate the second-order statistics of the error as the arithmetic average over a number of Monte Carlo realizations for every column vector $\mathbf{e}_l$ in all the K epochs. The resultant mean RMSE thus varies with the sampling index l as well as the index k. In order to obtain a scalar performance metric, we may average these errors over a certain time span. Recursive nonlinear estimators require a few sets of measurements, before converging to their stationary tracking performance. Therefore, we start averaging the values of the mean RMSE after an initial transient lasting two snapshot periods:

$$\overline{\text{RMSE}} = \frac{\sum_{l=2N+1}^{L} \text{RMSE}_l}{L-2N} \tag{3.53}$$

where

$$\text{RMSE}_l = \sqrt{\text{E}_{p\left(\mathbf{x}_k,[\mathbf{y}_1 \dots \mathbf{y}_k]\right)}[\mathbf{e}_l \mathbf{e}_l^T]}, \quad l \in \{(k-1)N+1, ..., Nk\} \tag{3.54}$$

In other words, the errors preceding the third observation are discarded to provide a metric less affected by the convergence time. If the time- and statistically-averaged RMSE attains the PCRB, the estimator employed is optimal: no better solution exists to minimize the MSE. This is the case of the standard KF for linear estimations problems.

3.2.3 Tracking Performance

The goal of the following simulation campaign is to investigate the impact of the rate of the snapshots carrying the TDOA/FDOA observables on the RMSE of the jammer position and velocity estimation. For this task, we test a recursive LS estimator and various KFs. These filters, namely the EKF1, the EKF2, and the UKF, differ in the order of approximation for propagating the information through the nonlinearities. The questions we are going to answer are: what should be the requirement on the snapshot period for a suitable tracking performance? What is the accuracy both outside and inside the area delimited by the receivers in this respect? Is it more convenient to rely on either long and sporadic measurements or short and frequent ones?

Simulation Scenario

The sensitivity of the estimation to the measurement precisions is magnified by the geometry of the arrangement of the DoOs with respect to the jammer. The GDOP is the ratio between the deviation of the position and velocity at the output of the estimator and the deviation of the input TDOA and FDOA observables. The higher it is, the worse is the quality of the transmitter-receiver geometry. By assuming that the measurements are uncorrelated and their noise covariance matrix is known, for example from Eq. 3.31, we may define a dilution of precision at the epoch k as a matrix

$$\mathbf{DOP}_k = \sqrt{\frac{\mathbf{PCRB}_k}{\mathbf{R}_k}} \tag{3.55}$$

according to [139]. This generic metric is neither scalar nor purely geometric, because it depends on the time (i.e., k) and also incorporates the memory effect of the dynamic model (i.e., $\mathbf{A}$ and $\mathbf{Q}$) through FIM in Eq. 3.45. Therefore, let us consider only the information related to the observation (i.e., $\mathbf{H}_k$ and $\mathbf{R}_k$) at the first time epoch (i.e., $k = 1$), hence when no recursion has been performed yet. In principle, this is the same as restricting ourselves to the first iteration of the original LS estimator. In this perspective, we may re-calculate the PCRB in Eq. 3.55 by initializing the estimation at any arbitrary guess $\hat{\mathbf{x}}_0$. The GDOP finally results from the sum of the diagonal elements as follows

$$\text{GDOP}_k = \sqrt{\text{trace}\left(\frac{(\mathbf{H}_k^T \mathbf{R}_k^{-1} \mathbf{H}_k)^{-1}}{\mathbf{R}_k}\right)} = \sqrt{\text{trace}\left((\mathbf{H}_k^T \mathbf{H}_k)^{-1}\right)}, \quad k = 1 \tag{3.56}$$

which we refer to from this moment on by omitting the index k. The precision of the measurements does not appear in the definition, so that the resultant metric is actually scalar, geometric, and independent of the time.

Since four are the components in $\mathbf{x}_k$, four are the minimum number of simultaneous observables $\mathbf{y}_k$, which are necessary to solve the equations that relate the position and velocity of the jammer to the TDOA/FDOA pairs. In the following, we rely so on solely four DoOs, which provide the minimum necessary number of measurements to resolve the four unknowns. These receivers are placed at the corners of a square with 100-m side. Changing their arrangement would only affect the GDOP, which for TDOAs has been already been widely investigated in the literature, also in the context of jammer localization (such as in [139]). The two-dimensional locations of the DoO in the scenario under study are marked in the geometry map of Fig. 3.6, which reports the GDOP for any hypothetical initial location $\hat{\mathbf{x}}_0$ of the transmitter in the area surrounding the receivers. The GDOP is unitary and thus ideal inside the square convex hull that is delimited by the DoOs, while it grows with the distance outside. It apparently tends to infinity at the locations that coincide with any of the receivers, as we have already noticed by looking at the elements in the $\mathbf{H}_k$ matrix. Such a proximity would also realistically cause the saturation of the front end, which is a possibility we do not take into account. With the aim of testing the performance for low and high GDOPs separately, we let the jamming attack originate at two possible locations: one inside and one outside the area surrounded by the DoOs. Since the only propagation effect on the signal is due to the isotropic free-space path loss and the devices are arranged in a square, the GDOP is symmetric in all the directions. This means that we do not need to evaluate the performance in tracking a jammer that moves across the whole scenario. We rather exploit this spatial symmetry of the attenuation to restrict the tests to smaller regions. The two regions chosen inside and outside the area are shown in Fig. 3.7, enclosed within the circles inscribed into triangular cuts of the scenario. The jammer starts moving from the incenters of these circles, which are located in (35.36 m, 14.64 m) and (-49.65 m, -49.85 m). It is headed in a straight and random direction at 100 km/h. The simulation time set is too short for the jammer to travel distance longer than the radius (i.e., about 14.64 m), so that it can never cross the borders of these regions.

The main simulation parameters are listed in Tab. 3.1. The sampling frequency f_s is set according to the specifications of common SDRs, while the observation time and the visible spectrum are wide enough to increase the input JNR (i.e., ρ) by a processing gain equal to 40 dB. Notably, the integration time is short enough to keep up with high-speed jammers far faster than the one tested. In fact, for any DoO in the scenario under test in Fig. 3.7, within this interval the maximum possible variations of delay and Doppler frequency shift may be roughly and respectively determined as

$$|y_\tau(t+T_{\text{obs}}) - y_\tau(t)| < 100 \ \ \text{km/h} \, \frac{T_{\text{obs}}}{\text{c}} = 0.09 \ \ \text{ns} < \frac{1}{f_s} = 50 \ \ \text{ns} \tag{3.57}$$

$$|y_\nu(t+T_{\text{obs}}) - y_\nu(t)| < \frac{f_c}{\text{c}} 100 \ \ \text{km/h} = 146.97 \ \ \text{Hz} < \frac{f_s}{N} = 1 \ \ \text{kHz} \tag{3.58}$$

which mean that the constraints in Eqs. 3.22 and 3.23 are fulfilled.

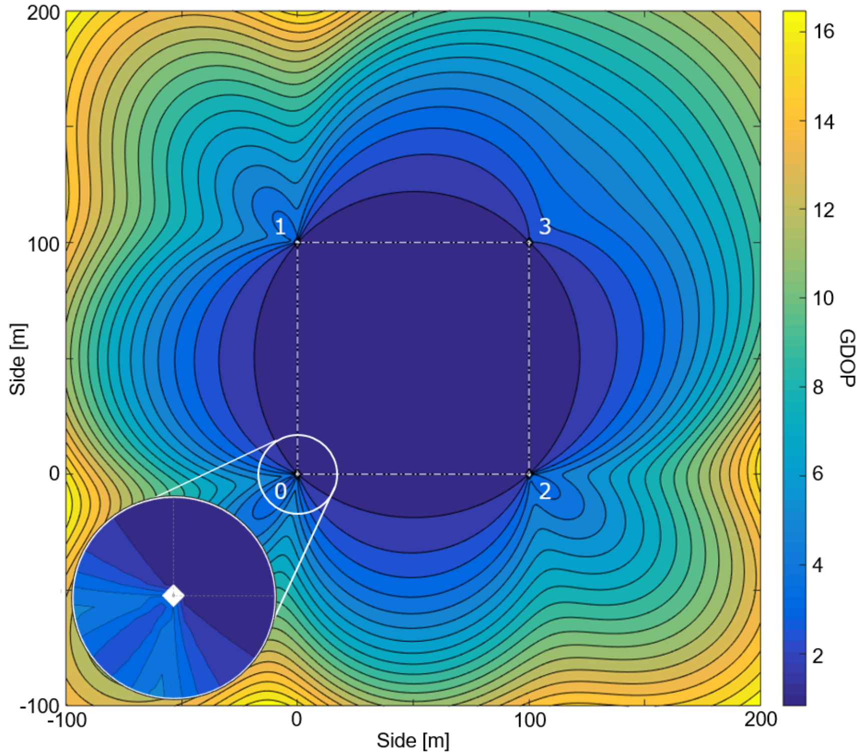

FIGURE 3.6: Map of the GDOP for the simulation scenario.

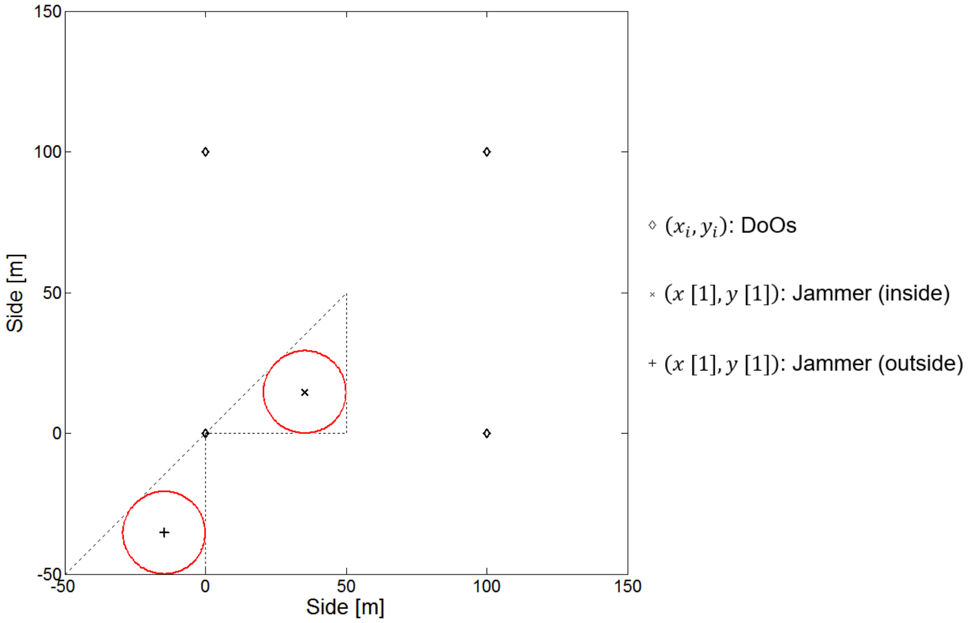
150
100
50
0
-50
Side [m]
-50
0
50
100
150
Side [m]
◇ (x_i, y_i): DoOs
× $(x[1], y[1])$: Jammer (inside)
+ $(x[1], y[1])$: Jammer (outside)

FIGURE 3.7: Inside and outside regions of the area between the DoOs that are tested for the jammer state in the simulation campaign.

Numerical Results

The first collection of simulation results characterizes the accuracy achieved by the ML estimator in retrieving the TDOA/FDOA pairs from the jamming signal. The interference is generated from (35.36 m, 14.64 m), namely the center of the circular

TABLE 3.1: Simulation parameters in tracking a single jammer.

Description	Symbol	Value
Carrier frequency	f_c	1575.42 MHz
Sampling frequency	f_s	20 MHz
Single-sided noise power spectral density	N_0	-194 dBW/Hz
Front-end equivalent noise passband bandwidth	B_{eq}	10 MHz
Jammer passband bandwidth	B_{RF}	10 MHz
Jammer repetition period	T_0	10 µs
Jammer transmit power	P_t	{-50, -40, -30, -20, -10, 0} dBW
Observation time	T_{obs}	1 ms
Normalized snapshot rate	R_s	{1, 1/5, 1/25, 1/125}
Snapshot period		{0.001, 0.005, 0.025, 0.125} s
Simulation time	T	0.5 s+T_{obs}
Number of observations	K	501
Number of snapshots		{501, 101, 21, 5}

region inside the area delimited by the DoOs. This test is meant to evaluate the observables that represent the outputs at the first step of the estimation process and the measurements at the input of the second step, which recursively estimates the jammer state. Such a characterization is carried out by averaging the RMSE separately for the TDOA and the FDOA over the initial observation (i.e., $k = 1$). In the following figures, the dashed lines always indicate the 95% confidence interval for the random errors returned by the simulations, under the assumption of a normal distribution. The outcomes are compared to the CRB in Eqs. 3.13-3.15, which are also later used to populate the matrix $\mathbf{R}_k$ in Eq. 3.31. The TDOA does not attain the lower bound because of the small bias that is introduced by the quantization of the delay-Doppler domain. The zero-padding interpolation of the FFT points in 3.20 refines the grid by shrinking the time spacing from 50 ns to 5 ns, which corresponds to 0.3 m. Evidently, the addition of points only partially compensates for the inevitable energy leakage in the discrete CAF. As a result, the delay accuracy in Fig. 3.8a saturates with the transmit power. The finer granularity is more effective in frequency, where the spacing between FDOAs is of 10 Hz, instead of 1 kHz, and the performance get relatively closer to the CRB in Fig. 3.8b. A small offset is still present at the high power levels, where the noise is negligible, again because of the quantized domain. Adding more zeros to the FFT could further enhance the precision TDOA/FDOA observations, until the error due to quantization losses is more significant than the resolution after interpolation. Anyway, we consider the results in Fig. 3.8 to be satisfactory, because higher gains in terms of accuracy might not be worth extra computations of longer FFTs. The CRB curves in these figures are valid, because the SNR at the output of the CAF far exceeds the 10 dB imposed by the conditions underlying Eqs. 3.14 and 3.15. This fact is evident from the input JNRs in Fig. 3.8c.

Now, we analyze the final performance of the two-step estimation process by evaluating the RMSEs for the position and the velocity. The jammer starts from the same location within the area monitored by the devices and the simulation terminates after about 0.5 s, when it has traveled by 13.89 m in any random direction. In the meanwhile, a time span equivalent to 501 observations has pasted and a variable

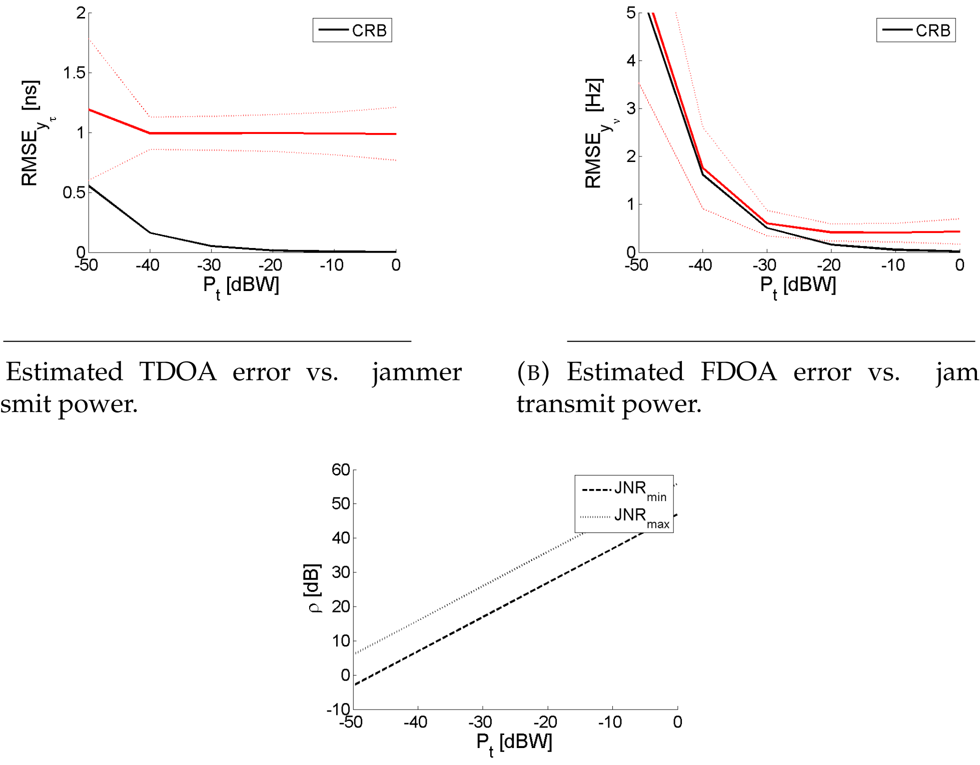

(A) Estimated TDOA error vs. jammer transmit power.

(B) Estimated FDOA error vs. jammer transmit power.

(C) Maximum and minimum JNRs at the four DoOs vs. jammer transmit power.

FIGURE 3.8: Performance of the joint TDOA/FDOA estimation at the first time epoch with the jammer moving from (35.36 m, 14.64 m).

number of snapshots has been taken. At the beginning, in (35.36 m, 14.64 m), all the recursive estimators are initialized with

$$\hat{\mathbf{x}}_0 = [50 \;\; \mathrm{m} \quad 50 \;\; \mathrm{m} \quad 0 \quad 0] \tag{3.59}$$

and, as far as the KFs are concerned, also

$$\hat{\mathbf{P}}_0 = \mathrm{diag}\left((50\,\mathrm{m})^2 \; (50\,\mathrm{m})^2 \; (200\,\mathrm{km/h})^2 \; (200\,\mathrm{km/h})^2\right). \tag{3.60}$$

where the values are chosen in view of the scenario, namely the size of the area (i.e., square with side 100 m) between the DoOs and the range of speeds expected for the jammer (i.e., at most 200 km/h). As for the model in Eq. 3.41, the state noise covariance is well described by the matrix

$$\mathbf{Q}_{\mathrm{low}} = \mathrm{diag}\left((0.1\,\mathrm{m})^2 \; (0.1\,\mathrm{m})^2 \; (1\,\mathrm{km/h})^2 \; (1\,\mathrm{km/h})^2\right) \tag{3.61}$$

the elements of which are set coherently with the previous characterization of the joint TDOA/FDOA estimation. The observation time is fixed to 1 ms, such that the simulation results in Fig. 3.9 are returned at equal processing gain, thus with the same measurement precision. They demonstrate how the KFs improve on the LS estimator, the location estimate of which degrades exponentially with respect to the snapshot rate. In fact, the accuracies of these filters do not deviate from a sub-meter position error, regardless of the snapshot period. This enhancement is possible because the constant-speed dynamic model is consistent with the scenario tested and,

consequently, the velocity estimation performs fairly well inside the area, where the GDOP is practically unitary. Therefore, the recursive prediction of the next state succeeds in correctly propagating the knowledge of the jammer state, despite the lack of observables. More in detail, even though it inherits the errors from both the joint TDOA/FDOA estimation and the approximation of the TDOA and FDOA nonlinear equations, the final bias achieved is about 0.4-0.5 m for the jammer position, and between 5 km/h and 10 km/h for the velocity, which is obviously the same for all the estimators used. As for the location estimation performance, the behavior of the UKF represents the exception, though, because it experiences an apparent degradation at low rates and high power. The position accuracy actually drops due to a slower convergence time. As one may notice from the example in Fig. 3.10, the initial transient indeed exceeds the first two snapshot periods that we have discarded in Eq. 3.53 to average the RMSE over time. In other words, three measurement updates are not sufficient for the UKF to converge, which requires a fourth snapshot before outperforming the LS estimator like the other KFs. Moreover, a well-known improvement for this filter is achievable by adding extra sigma points, which are used to propagate both the state and the measurement noise distributions (i.e., $\mathbf{w}_k$ and $\mathbf{v}_k$, respectively) through the nonlinearities. This upgrade is carried out by augmenting the vector and the covariance matrix of the state, both in prediction and update equations. Although this addition is straightforward, our filter includes only $\mathbf{w}_k$ and not $\mathbf{v}_k$ into the augmented state covariance, because this latter one produces an ill-conditioned matrix. Some attempts were made to overcome this issue by resorting to a more numerically robust implementation based on the square-root UKF in [155], but without success. This fact is arguably the reason why the performance of our filter does not differ from those of the extended KFs, since the augmentation is only applied to the linear dynamics model only and not the nonlinear observation.
Fig. 3.11 plots the results obtained at constant snapshot rate by increasing the duration of the observations and so the processing gain. This configuration increases the precision of the measurements to the detriment of availability of new observables, which are extracted from snapshots spread in time. Given the measurement precision already achievable with 1 ms of integration time, the LS estimator performs better when opting for frequent but short snapshots, rather than sporadic and long ones. On the contrary, and similarly to the previous results, the KFs exhibit marginal dependence on the observation time. While low state noise (e.g., in 3.61) is suitable to characterize the dynamics in case the jammer is moving inside the area between the receiving devices, high noise is necessary to reflect the higher uncertainty outside the area due progressive growth of the GDOP with distance. In the scenario under test, we model this latter case with

$$\mathbf{Q}_{\text{high}} = \text{diag}\left((1\text{ m})^2\ (1\text{ m})^2\ (10\text{ km/h})^2\ (10\text{ km/h})^2\right) \tag{3.62}$$

that we adopt to obtain the results shown in the Fig. 3.12. Although the observables are still precise, the poor GDOP poses the stability of KFs at risk: it can mislead the predictions based on the dynamics model, of which these filters take advantage in the absence of new measurements. The higher state noise in Eq. 3.62 is meant to preserve the convergence of the filters, but it also slows it down. For instance, the EKF2 has an unstable behavior. The EKF1 and UKF converge, but their improvement with respect to the LS estimator is reduced by the noisy dynamics model. The final biases are so generally on the orders of few meters for the position and tens of kilometers per hour for the velocity. The UKF exhibits the least residual dependence on the snapshot rate. From the results in the two regions of the scenario, both inside and

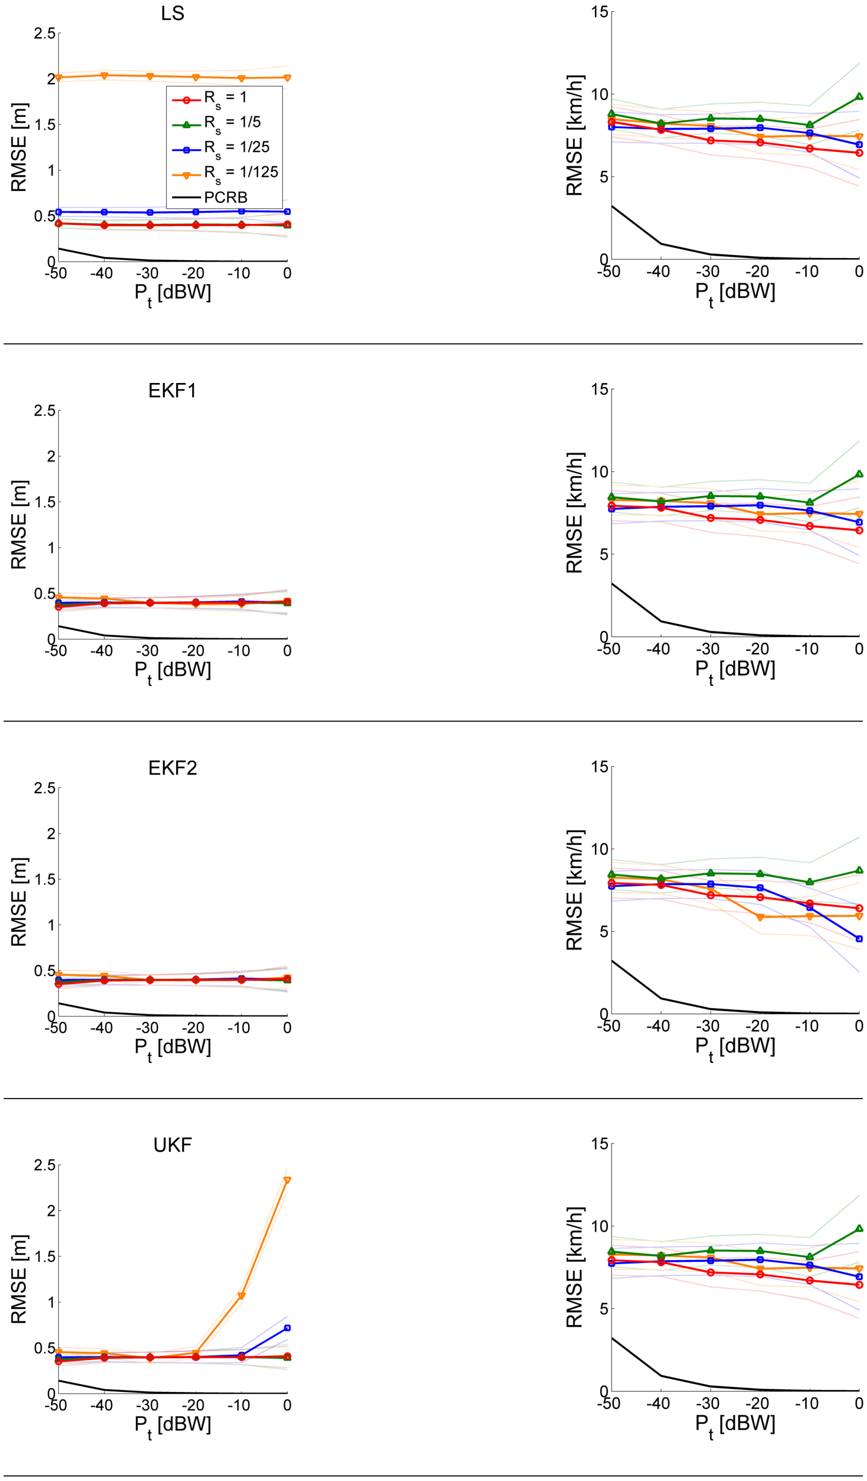

FIGURE 3.9: Average tracking performance vs. jammer transmit power at different snapshot rates inside the area between the DoOs.

outside the square area delimited by the DoOs, it is not possible to infer the impact of the snapshots on the speed estimation. For any filter, in fact, the percentile curves are too loose to make any relevant consideration about the average RMSEs.

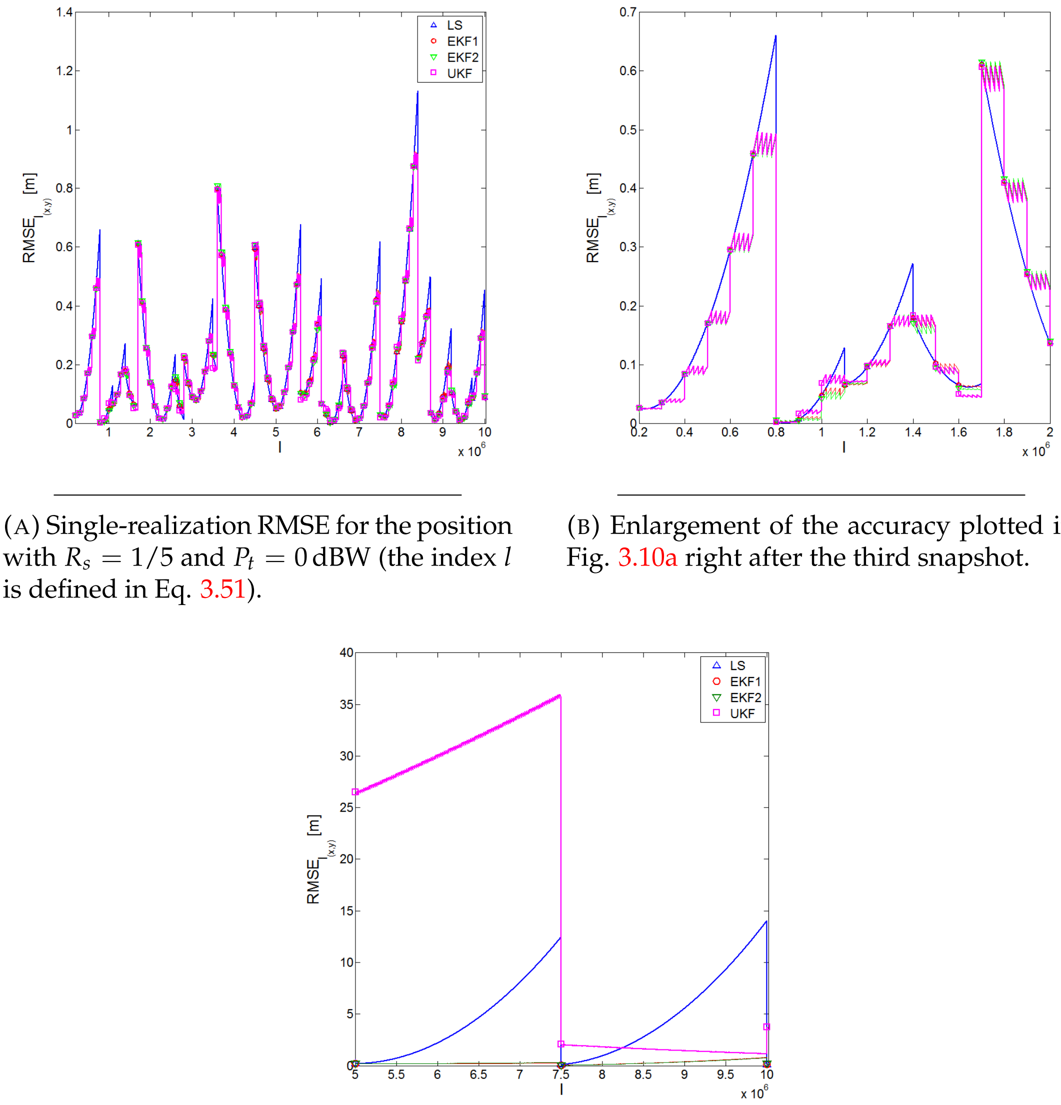

(A) Single-realization RMSE for the position with $R_s = 1/5$ and $P_t = 0$ dBW (the index l is defined in Eq. 3.51).

(B) Enlargement of the accuracy plotted in Fig. 3.10a right after the third snapshot.

(C) Single-realization RMSE for the position over time with $R_s = 1/125$ and $P_t = 0$ dBW.

FIGURE 3.10: Temporal evolution of the performance in tracking a jammer inside the area between the DoOs for a set of R_s.

3.2.4 Conclusions

Despite the intermittent availability of measurements, the jammer is localized with sub-meter accuracies where the GDOP is good, namely in the convex hull surrounded by receiving devices. In this area, a sufficiently fast pace of snapshots is supposed to enable the continuous and accurate tracking in time of both position and velocity, by means of well-known variants of the KF. Such a filter works with a simple dynamics model that assumes a stationary trajectory at constant speed. Whenever the model well approximates the jammer motion between two consecutive snapshots, the prediction functionality of the KF overcomes the temporary absence of TDOA/FDOA pairs. This capability facilitates the usage of sensing devices equipped with basic front ends for capturing snapshots of the jamming signal, because it allows for relaxing the constraints on their computational and energy resources. These sensors are used just to record samples of interference, because the processing load is dumped

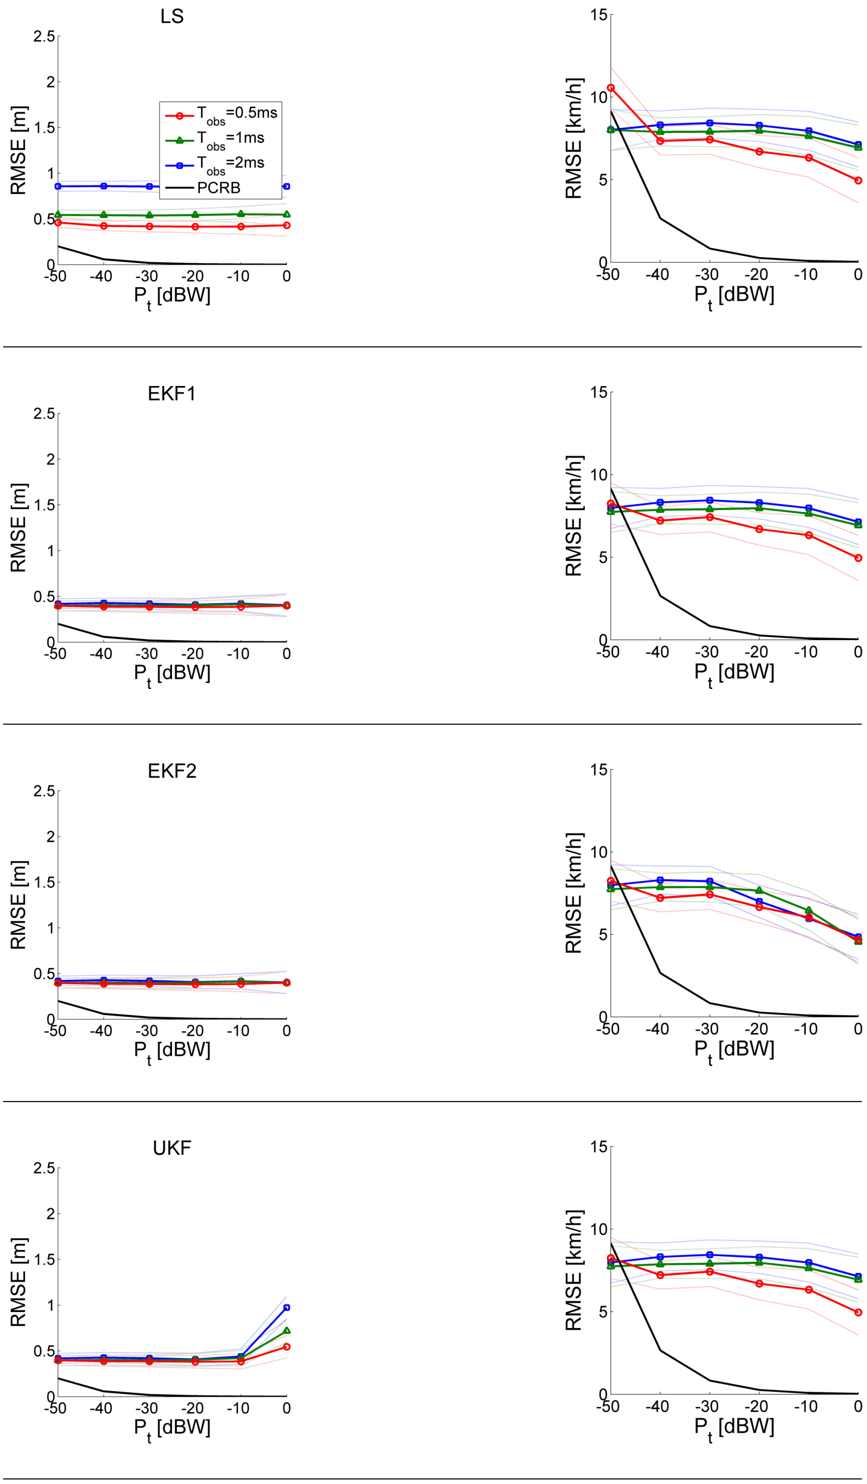

FIGURE 3.11: Average tracking performance vs. jammer transmit power for a set of T_{obs} inside the area between the DoOs.

to a central processor, which recursively estimates the current jammer location and speed in a timely manner. Whenever snapshots are made available by the receivers, the additional computation of the CAF is carried out through the low-complexity FFT, for the fast extraction of the observables. While the simulation campaign only considers a worst case with the minimum of four sensors, realistic application scenarios benefit from the deployment of many more receivers, which likely boost the

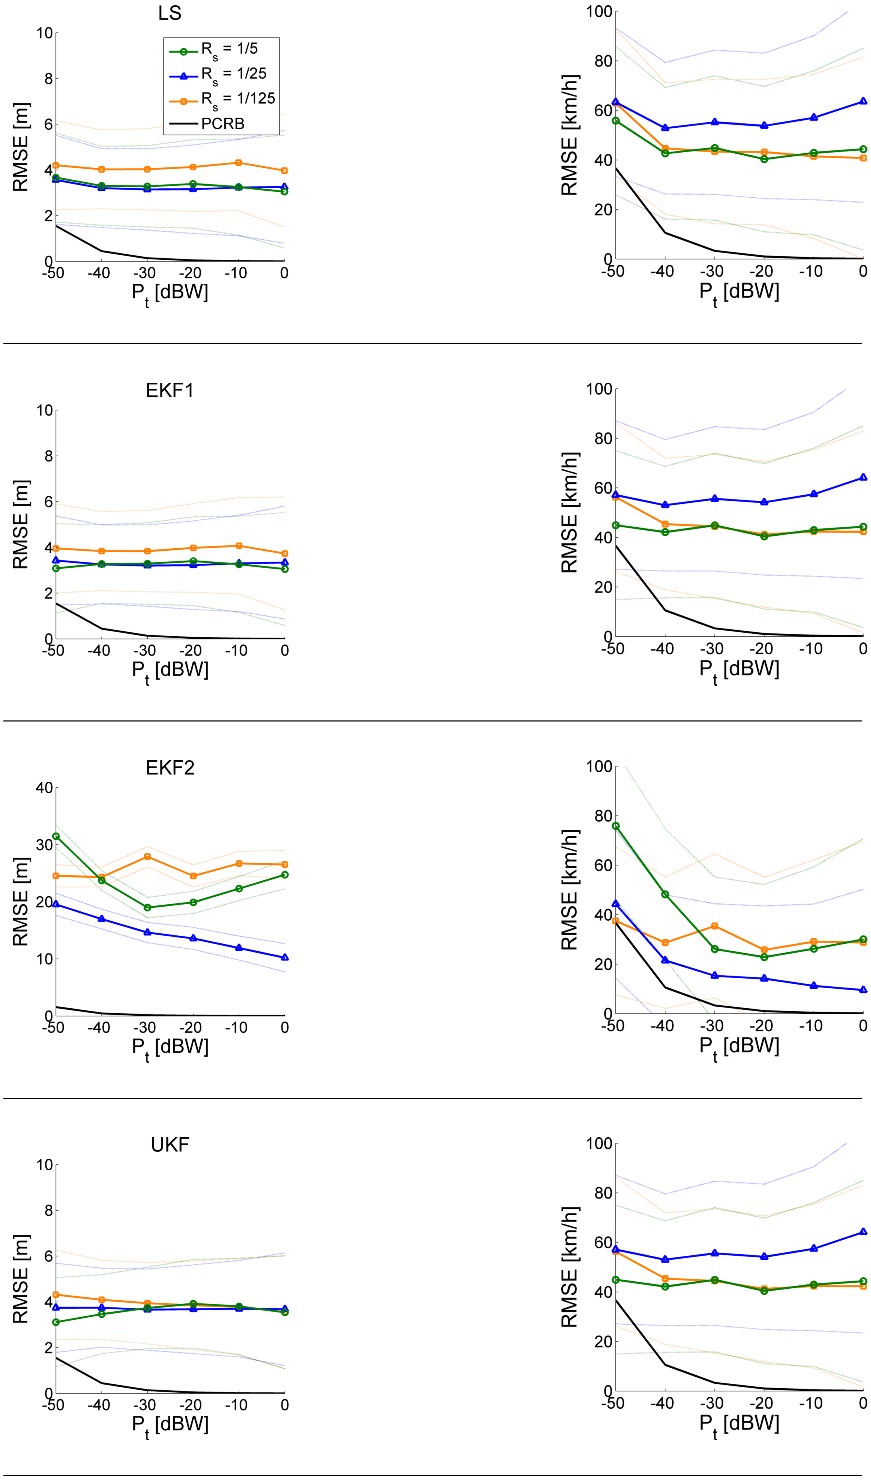

FIGURE 3.12: Average tracking performance vs. jammer transmit power for a set of R_s outside the area between the DoOs.

geometry and the redundancy of the observables. Moreover, the region that is monitored with the desired accuracy can be enlarged by distancing the devices, as long as these latter ones are evenly distributed in the area. Under this circumstance, indeed, a good GDOP is ensured without compromising on the precision of the measurements, which are weakly sensitive to the range between the transmitter and the receivers. All these aspects are discussed through the previous subsections and give important insights about the feasibility and the potential behind the crowdsourcing of TDOA and FDOA observables for jamming localization. In this perspective, we may conceive a system to aid law enforcements against sources of interference that exploits low-cost and low-power DoOs on a large scale and the processing power nowadays offered by the cloud. This application is ultimately one that could emerge under the forthcoming paradigm of the IoT, in which the information accessed by spatially-distributed sensors is merged for various purposes.

3.3 Characterization of Multiple In-Car Jammers

In the previous section, we have seen the two challenges in retrieving through cross-correlations the TDOA/FDOA pairs, when more than one jamming waveform are received. The first issue is that the ML estimator is not reliable, because the underlying model is not valid anymore: the overall CAF results from the combination of functions that are generated by the different waveforms and not one. Thence, if the jammer are closely-spaced and/or moving in the same direction, their main lobes might overlap in the delay-Doppler domain. When this chance occurs, a constructive interaction among the CAFs produces one single peak of magnitude, which causes the missed detection of one or more sources of interference. This situation is exemplified in Fig. 3.4. Likewise, the lobes may partially or totally cancel each other if the interaction is destructive, thus increasing the probabilities of missed detections as well as false alarm. For instance, the coherent summation of spurious correlations in the CAFs can give rise to significant peaks, which are nothing but artifacts that likely lead to false alarms. Usually, jamming waveforms comply with the model underlying Eq. 3.24: their magnitudes have a sinc-like shape featuring side lobes, as shown in Fig. 3.3. All in all, finding the peaks of energy in the delay-Doppler domain is not a suitable way of resolving the superposition of the waveforms received. Even assuming that the joint estimation TDOA and FDOA is correctly accomplished, the second issue is then identifying which of the observables extracted from every couple of synchronous receivers belong to a certain jammer, given that the number of sources is unknown. This challenge is referred to as the data association problem and is generally solved through complicated iterative algorithms iteratively. In [140], the estimation process for localizing multiple jammers solves the first issue as an optimization problem that is built around the model in 3.24. The associations are then overcome by clustering the resultant TDOA/FDOA pairs in the position domain. Alternatively to this approach, the strategy we propose in the following is meant to bypass the two aforementioned challenges through TF analysis. Particularly, we leverage on the knowledge about the commercial jammers that are typically on board of civilian vehicle in order to recognize and separate the signatures of the jamming waveforms embedded into the incoming signal. Thanks to field investigations and experimental surveys, such as [146], we known that most of in-car PPDs transmit linear chirps that are modulated in frequency according to a sawtooth pattern. We may take advantage of these notions to distinguish the concurrent jamming attacks based on the differences between the respective TF characteristics. In the present subsection, we review an algorithm that is conceived to address this task. More specifically, this algorithm provides "one-shot" estimates of the repetition periods and the rates of the sawtooths that overlap in time and frequency. Clustering these estimates in a sort of TF domain refines the characterization and counts the number of simultaneous jammers. The extension to a tracking operation similar to that in [131] for one jammer is possible but not yet explored. The capability enabled by this algorithm may turn out to be useful in countermeasures against interference. Later in this chapter, it is indeed used to inject additional information into the two-step estimation process that can track the positions and velocities of the two jammers.

A *single* receiver is being employed to infer the TF characteristics of an *unknown* number of *multiple* jamming waveforms, which are superimposed with arbitrary phases and time-varying amplitudes. They are *periodic*, highly *non-stationary* (i.e., their spectra change rapidly), and generated by jammers with various transmit powers and/or at diverse distances. As before, the possible reception of multipath replica is not taken into account for the sake of simplicity. The useful signals victim of interference are assumed to have a power level well below the noise floor (e.g., DS-SS communications and GNSSs), so that they do not affect the characterization of the jammers. On the contrary, the power spectra related to the jamming attempts are supposed to exceed the noise power density within the visible spectrum. Furthermore, the receiver should be *continuously* jammed over the time of observation. This last hypothesis excludes the possibility of pulsed behaviors during the reception and simplifies the algorithm. Nevertheless, in practice, repetitive observations could be combined to relax this assumption. The goal of the application under study is determining the number and the respective inner periodicities of the simultaneous waveforms observed at the receiver end. Thereupon, the incoming mixture of energy components is analyzed onto a two-dimensional TF domain, rather than one spectral dimension in order to discriminate among their frequency modulations. The TF analysis is performed by digitally processing the signal received, without the need of additional analog components at the front end. The ST is an effective tool to tackle the uncertainty about the overlapping modulations, because of its desirable properties introduced in Chapter 2. This transform is linear and capitalizes on the insertion of Gaussian windows with deviations progressive in frequency into the exponential basis functions that underlie the conventional Fourier analysis. By doing so, it can provide consistent multi-resolution TF representations that are free from fictitious cross terms and that identify the local timings of both amplitude and phase components, which are otherwise averaged over time into the global FT. Therefore, the algorithm explained in next subsection searches for jamming waveforms by computing the discrete and complex-valued ST for *snapshots* of received I/Q samples in a *batch* fashion. Despite the high complexity of the ST, a fast computation is still feasible if the TF resolution is down-scaled by limiting the snapshot to short time spans (i.e., short observation times) and/or the sampling rate. Since the snapshot are long enough to provide sufficiently high time-bandwidth product and so processing gain, while being short enough to be processed fast through the fully-redundant ST, the implementation of the temporal decimation and spectral compression described in Chapter 2 are not investigated for this application; reducing the redundancy to limit the complexity is unnecessary and inevitably introduce some errors. In any case, the receiver can be simply used as a basic front end, while the snapshots are analyzed on the *cloud*, which is endowed with enough processing power to perform jamming characterization *online*, without compromising on the resolution. Consecutive snapshots are hereinafter treated as independent realizations of the signal, because the number of sources in the scenario might change in the meanwhile and no tracking operation is considered. Some additional capabilities provided by the algorithm presented, and yet not evaluated in the following, are interference detection and the characterization of stationary in-band disturbances (e.g., continuous waves). The search for jamming waveforms is supposedly triggered only if a relevant amount of energy exceeds the noise floor in the ST for the current snapshot. This mechanism not only avoids to waste computational resources for extracting periodicities from empty batches, but also flags the presence of strong interference. The identification of any waveform not modulated in frequency, which is possibly generated by malfunctioning equipments or powerful co-existing systems (e.g., radars), is straightforward in the ST. The algorithms for the characterization of jamming waveforms is robust to the reception of these additional disturbances.

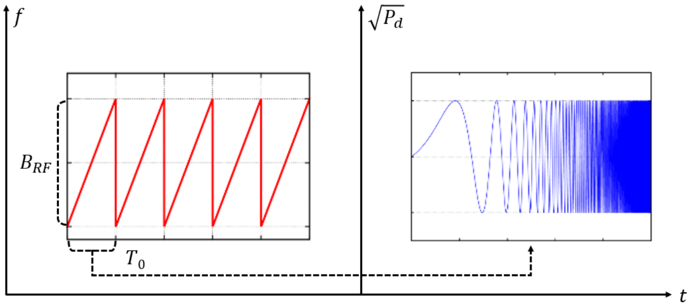

FIGURE 3.13: Sawtooth pattern modulating in frequency a single jamming waveform generated by in-car jammers.

TABLE 3.2: Typical characteristics for the waveforms of the transmitters jamming GPS bands.

Symbol	Range	Average
B_{RF}	10-60 MHz	15 MHz
T_0	1-30 µs	10 µs

3.3.1 Optimization of the ST for Jamming Sawtooth Waveforms

Let us resume the notation in Eq. 3.1. After the down-conversion to baseband, the complex envelope of the analog signal consists of a number N_ψ of jamming waveforms superimposed in time and frequency and defined as

$$d(t) = \sum_{i=1}^{N_\psi} \left(\sqrt{P_{d_i}(t)} \sum_{r=-\infty}^{\infty} \psi_i\big(t-(r-1)T_{0_i}\big) \mathrm{e}^{j\phi_i(t)} \right) \tag{3.63}$$

where $\psi_i(t)$ is the sawtooth function that modulates the i-th train of linear chirps according to Eq. 3.2, and $\phi_i(t)$ is an arbitrary phase shift due to a certain delay and Doppler frequency shift. The periodic repetition of chirps describes a pattern depicted in Fig. 3.13 and acts as the TF signature for the emission of the respective source of interference. According to the scientific literature, the vast majority of in-car jammers on the market are banks of transmitters that radiate interference onto distinct sub-bands through monopole antennas of suitable sizes. Every transmitter generates one sawtooth waveform that rapidly scans back and forth the dedicated sub-band with a narrowband and powerful tone. This frequency modulation is roughly linear and provides a twofold advantage from the point of view of the malicious user. It concentrates the energy in time, while covering large bands in short intervals. As a matter of fact, the interference captured at the receiver end is wideband, even though the emission at the origin is delivered with limited power consumption. This smart design reduces the bulk and lengthens the battery life of in-car jammers, which may be then commercialized as inexpensive and portable devices. We refer as characteristics of the jamming waveform to those of the one transmitter that radiates in the sub-band of interest. In this context, common values of repetition period and bandwidths for the interference within GPS bands are surveyed with commercial devices in [146] and reported here in Tab. 3.2.
After the front-end filter and the quantizer, the jamming waveform is embedded into

the noisy and complex-valued digital signal $s[l]$ that is already defined in Eqs. 3.3 and 3.16. We analyze the time-varying spectral content of a snapshot made of N of these I/Q samples by means the discrete ST. The analysis, however, is not always fruitful. Despite the convenient properties, the ST cannot provide the same performance in terms of energy concentration regardless of the input set of sawtooths. Indeed, even inside the bandwidth free from side-effects (i.e., self-aliasing and time misrepresentation), the energy distribution of a frequency-modulated waveform is gradually spread the further the spectral components of the signal are away from the central frequency. This outcome could harm the correct interpretation of the resultant TF representation, to the point that the output of the ST is practically useless. The consequence is notably evident for fast modulations, in which several repetitions fall into the snapshot under analysis. This shortcoming comes from the progressive trade-off of TF resolution as originally formulated by [23]. In here, the deviation of the Gaussian function that windows the signals is linearly scaled so that it is equal to the reciprocal of the magnitude of the corresponding normalized frequency. This formulation lacks of flexibility: no matter what signal is analyzed, the windows become flatter as the frequency modulus is higher. The simple introduction of a variable coefficient of proportionality k in the Gaussian function provides an additional degree of freedom to control the resolution trade-off. This addition is suggested in [156] to improve the performance of the ST, which is constructed with the following windows

$$w(k)[n,p] = \frac{|p|}{kN\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{n}{k\sigma}\right)^2\right), \quad \sigma = \frac{N}{|p|} \tag{3.64}$$

for digital frequency bins identified by the index $p = -N/2, ..., N/2-1$ with $N = f_s T_{\text{obs}}$. Ergo, the ST defined in Chapter 2 is modified into the following matrix

$$\text{ST}_s(k)[n,p] = \sum_{m=l}^{l+N-1} s[m]w(k)[m-n]\text{e}^{-j2\pi\frac{p}{N}m} \tag{3.65}$$

and the condition on the sampling rate for minimizing the self-aliasing effect changes accordingly:

$$f_s \geq \frac{B_{\text{eq}}(k\pi+1)}{k\pi}. \tag{3.66}$$

The bound imposed by Eq. 3.66 is evidently much stricter than that of the Nyquist sampling criterion. Therefore, from this moment on, we sample the signal with a rate double with respect to the front-end passband bandwidth (i.e., $f_s = 2B_{\text{eq}}$) in order to grant the freedom to sweep k up to a maximum value equal to π^{-1} whether necessary. The number of oscillations that can be contained within one standard deviation of a window grows with the product $k\sigma$. While the temporal resolution benefits from a small coefficient, the spectral resolution is compromised, especially for low frequency magnitudes $|p|$. Conversely, if k is large, more cycles of the sinusoids are retained within one window and thus the frequency resolution is enhanced to the detriment of that in time, especially when $|p|$ is high. The complete control of the TF trade-off is gained with the adoption of a scaling factor that is a function with respect to the frequency (i.e., $k[p]$). One can strain or stretch the progression that regulates the effective duration of the Gaussian windows, by tuning this factor as needed. Ideally, the scaling function is tailored to the energy spectral density of the signal analyzed. For instance, a constant scale step is suitable to represent linear chirps, whereas a frequency-varying step is better for nonlinear waveforms. The adaptation of $k[p]$ requires the knowledge of the incoming signal. In [157] and the articles cited therein, the standard deviation is optimized *a posteriori* to maximize a measure of the energy concentration for every spectral component. These methods are computationally consuming and, most of all, are possibly inconvenient for our application, because the TF characteristics of the multiple jamming waveforms are diverse and unknown. For the sake of practicability, the *a-priori* optimization of this factor within the receiver bandwidth is feasible by making reasonable assumptions about the class of waveforms of interest. The class we consider is the one described as the summation of linear frequency-modulated waveforms in Eq. 3.63. Interestingly, as far as linear chirps are concerned, a closed-form approximation of the optimized standard deviation for fixed Gaussian window of the STFT is found in [41] based on one specific chirp rate. We may apply the same principle to the frequency-dependent windows of the ST and derive the formula:

$$k_{\mathrm{opt}} = 2\sqrt{\frac{\pi B_{\mathrm{RF}}}{2T_0 f_s^2}} \tag{3.67}$$

the demonstration of which is straightforward. Equivalently, this equation may be expressed with respect to the chirp angle θ as follows

$$k_{\mathrm{opt}} = 2\sqrt{\frac{\pi \tan(\pi - \theta)}{2N}} \tag{3.68}$$

over the grid of $N \times N$ points that identify the TF components in the ST matrix, as shown in Fig. 3.14. Obviously, neither the rate nor the angle of the chirp are known. Moreover, Eq. 3.68 optimizes the TF representation for a single sawtooth waveform only, while we are dealing with many overlapping ones. Sophisticated solutions to this problem resort to iterative a-posteriori adaptations that make $k_{\mathrm{opt}}(n, p)$ vary in time and/or frequency. They entail a significant computational burden and are beyond the scope of the application subject of this chapter. Therefore, it is preferable a workaround able to work a priori, without further pre-processing. A simple effective solution in this sense is fixing the scaling function to one value denoted by $\bar{k}_{\mathrm{opt}}$, which is optimized to the *average* slope in the range of possible chirp rates. We may circumscribe this range by making a pair of considerations, also illustrated in Fig. 3.15.

- The ST should accurately represent any sawtooth waveform that has two or more repetitions lying within the observation time, regardless of the bandwidth of the front-end filer. This condition ties the chirp angle to the maximum periodicity that should be visible in the representation, namely

$$T_{0_{\max}} = \frac{T_{\mathrm{obs}}}{2}. \tag{3.69}$$

  From the geometry of the TF domain, it is plain to obtain a higher bound in the form of $\theta \leq 135°$.

- If the slope of the sawtooth is slow over time, the sampling rate is sufficient to follow the evolution of the waveform without energy leakages. If the modulation is fast and thus the slope is high, the spectral energy of the waveform could leak among the points that underlie the quantized TF domain of the discrete ST, similarly to what happens in the CAF. As a consequence, the energy of these components is spread among the neighbors and possibly invisible in

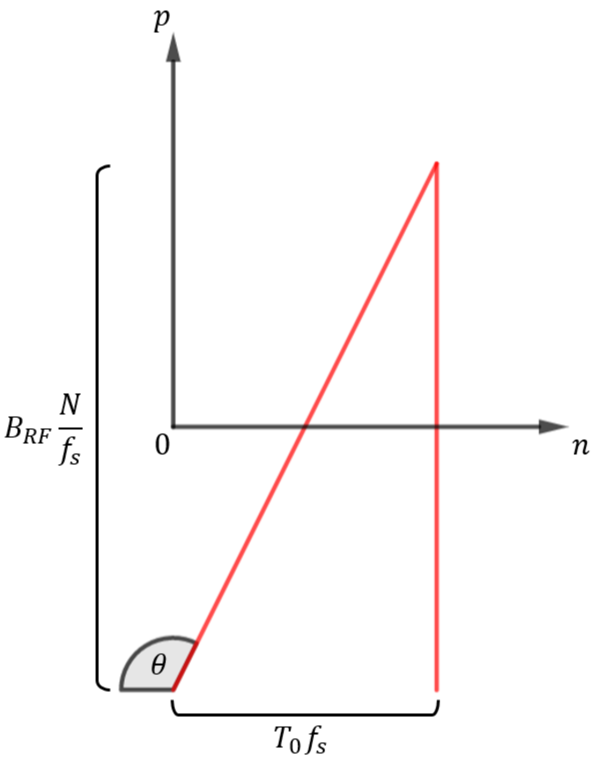

FIGURE 3.14: Definition of the chirp angle of a sawtooth waveform in the quantized TF domain under the discrete ST.

the resultant representation. To avoid this possibility, we set lower bound to the chirp angle with $\theta \geq 95°$, which also corresponds to a minimum repetition period:

$$T_{0_{\min}} = \frac{B_{\mathrm{eq}} N}{\tan(85°) f_s^2} \tag{3.70}$$

that depends on the receiver bandwidth. The threshold used to express the condition is not derived in a rigorous manner, but it is actually found by trial and error. The arbitrariness behind this setting is anyhow trivial, because the limitation it imposes can be overcome by adjusting observation time, as we will see in the results.

In the light of these considerations, the ST can be optimized with

$$k[p] = \bar{k}_{\mathrm{opt}}, \quad \forall p \tag{3.71}$$

according to

$$\bar{k}_{\mathrm{opt}} = 2\sqrt{\frac{\tan(\pi - \bar{\theta})}{2N}} \tag{3.72}$$

where the arithmetic average of the chirp angles is $\bar{\theta} = 115°$ within the range defined by $\theta \in [95°, 135°]$. This formula does not require any insight about the specific TF characteristics of the sawtooth waveform under analysis. The effectiveness of such a solution is demonstrated for example in Fig. 3.16.

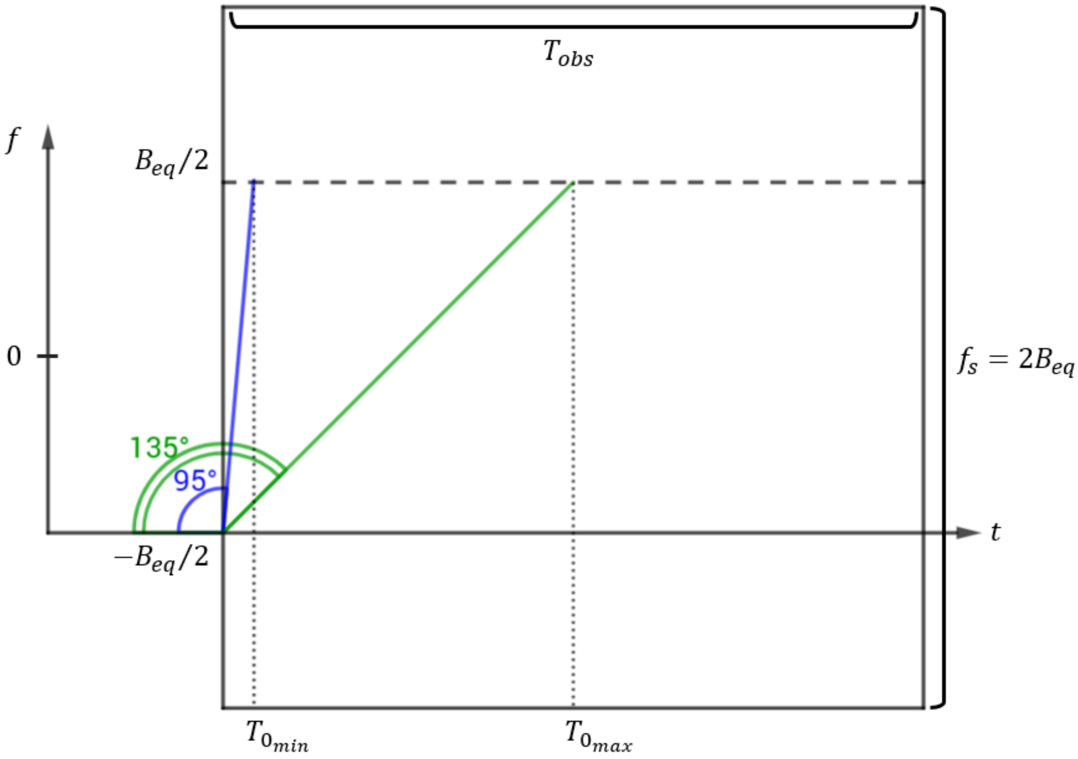

FIGURE 3.15: Chirp angle of the sawtooth waveform.

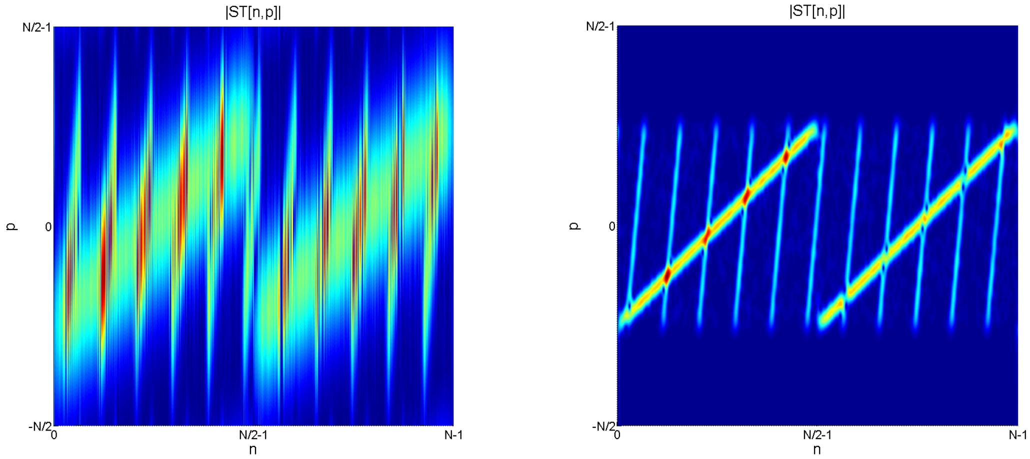

(A) Amplitude of the ST with $k[p] = 1$.

(B) Amplitude of the ST with $k[p] = \bar{k}_{\text{opt}}$.

FIGURE 3.16: Discrete TF representations over $T_{\text{obs}} = 100$ µs of two jamming sawtooth waveforms equally-powerful that are characterized by chirp angles $\theta = \{95.14°, 135°\}$ and have repetition periods $T_0 = \{10\ \mu\text{s}, 50\ \mu\text{s}\}$.

3.3.2 Energy Detection based on the S-Transform

Once being computed according to Eq. 3.72, the ST is optimized to analyze the sawtooth waveforms that jam the snapshots received. Although the resultant TF representation of the interference components is neat, the amplitudes and phases are of little use because they result from the untraceable constructive/destructive interactions among the waveforms. Therefore, the information they carry is not meaningful if the task is to separate and characterize the signatures of multiple jammers that are mixed into the TF domain. The idea is then distinguishing the sawtooth waveform by taking advantage of their inner periodicities. In other words, what matters is the recognition of the periodic evolutions in time of the energy distribution, rather than reading the corresponding values of amplitude and phase. To extract this information, the significant energy density of the interference should be first revealed. Any other intentional or unintentional emission within the receiver bandwidth is supposed to be weak enough to be buried under the noise floor. Thence, discriminating the interference from the noise is just a matter of detecting the TF components standing out in terms of energy. In this regard, within the spectrogram built on the STFT, a simple binary decision test is used in [43]. The detection statistics relies on the assumption of additive white Gaussian noise, the energy of which is proven to have a chi-squared PDF. Since the nature of this distribution does not depend on the analysis window, the same assertion holds in the framework of the ST. Under the hypothesis of the absence of signals, the probability for any TF component to exceed a certain energy threshold is

$$P_{\text{fa}}[p] = \text{prob}\left(|\text{ST}[n,p]|^2 > \lambda[p]\right) = \exp\left(-\frac{\lambda[p]}{2\sigma_n^2 E_w[p]}\right) \tag{3.73}$$

that is referred to as false-alarm probability, in which is the energy E_w of the frequency-dependent windows is known by the construction of the transform, whereas the knowledge of the noise power σ_n^2 is acquired during an initialization phase far from sources of interference. The threshold denoted by $\lambda[p]$ is a function with respect to the frequency because of the set of windows of the ST, while it has unique value for the STFT. Usually, the false-alarm probability is fixed by the following rule of thumb

$$P_{\text{fa}}[p] = \overline{P_{\text{fa}}} < \frac{1}{N^2}, \quad \forall p. \tag{3.74}$$

Therefore, the threshold can be pre-determined frequency by frequency as

$$\lambda[p] = -2\sigma_n^2 E_w[p] \ln\left(\overline{P_{\text{fa}}}\right). \tag{3.75}$$

Any component carrying an energy level above this threshold is flagged as jamming power. This statistics detects and locates the presence of interference in the TF domain. The outcome is a binary and two-dimensional mask, which is set to one in the correspondence to correct detections and false alarms, while it is null for the negligible TF components. The mathematical formulation of this concept is the matrix

$$\mathbf{\Lambda} = \Lambda[n,p] = \left|\mathbf{1}_N - u\left(\begin{bmatrix}\lambda[-N/2] & \dots & \lambda[N/2-1] \\ \vdots & \ddots & \vdots \\ \lambda[-N/2] & \dots & \lambda[N/2-1]\end{bmatrix} - \text{ST[n,p]}\right)\right| \tag{3.76}$$

where $\mathbf{1}_N$ is the all-ones square matrix of order N and u is the element-wise unit step function defined by

$$u(x) = \begin{cases} 1, & x \geq 0 \\ 0, & x < 0. \end{cases} \tag{3.77}$$

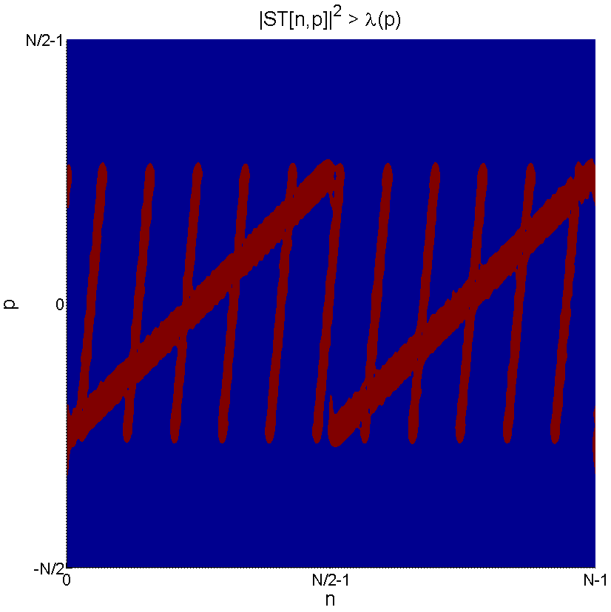

FIGURE 3.17: Binary mask computed from the ST in Fig. 3.15.

For example, the one based on the signal in the precedent figures is plotted in Fig. 3.17. We may notice that the energy of the fast modulation features thinner traces in the TF plane, while, on the contrary, the one with a slow rate exhibits a signature more spread in frequency, even though their power level are identical. Multiplying this mask to the matrix of the discrete ST means blanking all the information associated to the noise.

3.3.3 Estimation and Aggregation of the Sawtooth Periodicities

Since the interference generated by in-car jammer typically features a sawtooth pattern, we may regard periodicity as a further dimension to combine time and frequency information. And we can exploit this observable to distinguish the waveforms and to guess the respective TF characteristics, which are the chirp slope and the repetition time. In other words, we leverage on the little knowledge we have about the generic properties of the class of interference that we target in order to characterize the specific ones present in the scenario. Sophisticated techniques in the realm of advanced image processing and machine learning are believed to be promising when it comes to the best accomplishment of such a task. However, they are also too complex for the ultimate goal of this application, which is the tracking of simultaneous jammers discussed in the next subsection. It is then important to contain the execution time, so that the jamming characterization can run in nearly real time. For this purpose, the algorithm described below is devised to be as light as possible in terms of computations. The complexity could be handled by the computational resources at disposal on the cloud.
The binary mask defined by Eq. 3.76 discriminates the TF components of pure noise from those apparently affected by in-band interference. The distinction entails an implicit and reasonable choice: the search is restricted only to the waveforms that are received with enough power to exceed the background noise. Therefore, the receiver should be within the coverage of the jammers that have to be characterized,

as we consider hereinafter. The mask exposes the energy contained in the snapshot, regardless of the amplitude and phase at the reception. This information is then used as input into the algorithm that recognizes the signatures of the various jamming waveforms. The algorithm within a loop of two stages, which detect and estimate the periodicities in the mask, respectively. At the end of the loop, a third stage of clustering is in charge of aggregating the estimates to establish the number of jammers and retrieve both the respective chirp angles and repetition periods. The flowchart in Fig. 3.23 provides an overview of the whole operation. The three stages are summarized in the following.

Chirp recognition - Inside a loop, the coordinate system of the binary mask is rotated by an incremental angle, which is swept along a set drawn from the range of all possible chirp slopes defined in Fig. 3.15, namely from $\theta \in [95°, 135°]$. For the moment, let us consider only the set integer angles of $\theta_r = \{95°, ..., 135°\}$, so that the rotated matrix is denoted by $\boldsymbol{\Lambda}_r$ (or $\Lambda_r[n', p']$) for $r = 1, ..., R$ with $R = 41$ and a resolution of one degree. At every instance of the loop, the rotation is performed by multiplying the mask with a proper matrix that can be stored. The product is then inspected to reveal possible sawtooths characterized by the corresponding chirp slope. The set could be limited to a few specific angles, if one has already clues about the jammers present. After the rotation, the first stage of the search for jamming waveforms is based on the Hough transform (HT), which is a well-known tool to identify segments in images. This transform is introduced in [158] as a special case of the Radon transform that has a simple closed form. We use it to recognize line segments due to linear chirps into the binary mask. Processing the HT fixed at -90° for the matrix $\boldsymbol{\Lambda}_r$ returns a *one-dimensional* function. The task is to identify the *horizontal* and *continuous* sequences of TF components with significant energy (i.e., ones in $\boldsymbol{\Lambda}_r$) that exceed a certain length. The recognition finds at most a number N_{H} of bins of index p' where this function is higher than $B_{\mathrm{H}}(\theta_r)$ (see Fig. 3.19). Next, these bins are examined in order to extract sequences of contiguous ones, which should be at least longer than $B_{\mathrm{H}}(\theta_r)$ with the exceptions of gaps shorter or equal to M_{H}. This last parameter provides robustness against negligible discontinuities. The ultimate task is to reveal the directions along which the energy is distributed with continuity, thence the slope of the repeated chirps. An example is shown in Fig. 3.18. Once rotated, the axes of time and frequency coordinates have indices respectively equal to

$$n' = (n - \frac{N}{2}) \cos(\theta_r) - p \sin(\theta_r) + \frac{N}{2} \tag{3.78}$$

and

$$p' = (n - \frac{N}{2}) \sin(\theta_r) + p \cos(\theta_r). \tag{3.79}$$

The parameters N_{H} and M_{H} are arbitrary and have an important impact on the performance and the computational burden. If we restrict the algorithms to extract only the sawtooth that span the whole visible spectrum, we may tie the minimum length of the sequences extracted to the receiver bandwidth as follows

$$B_{\mathrm{H}}(\theta_r) = \lfloor \frac{0.9 B_{\mathrm{eq}}}{\sin(\pi - \theta_r)} \frac{N}{f_s} \rfloor \tag{3.80}$$

for obvious geometrical reasons.

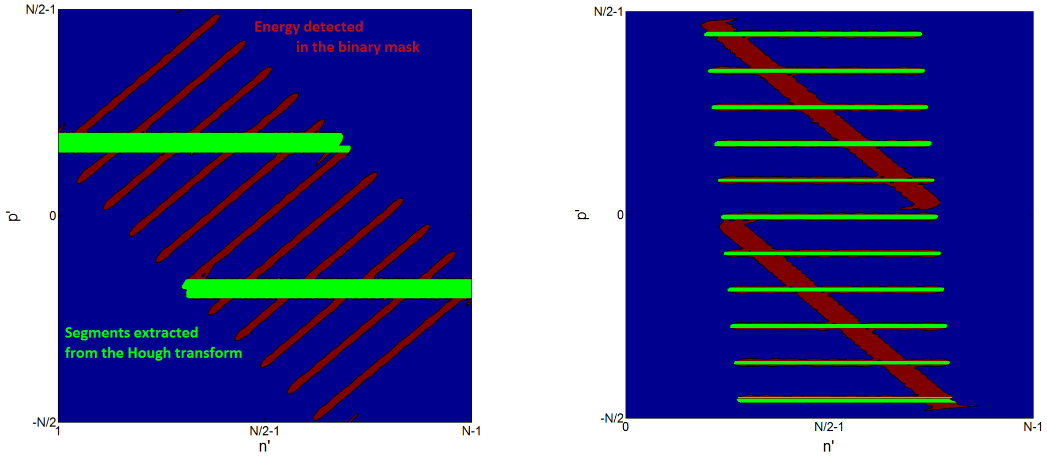

(A) Sequences extracted from $\mathbf{\Lambda}_{135°}$.

(B) Sequences extracted from $\mathbf{\Lambda}_{95°}$.

FIGURE 3.18: Continuous energy recognized in the mask of Fig. 3.17 for the two rotation angles $\theta_r = \{95°,135\ °\}$ that roughly coincide with slopes of the two jammer slopes.

Periodicity estimation - After being identified with the HT, the relevant concentration of energy in the matrix $\mathbf{\Lambda}_r$ is summed up horizontally, row by row, thus over the rotated time axis with index n'. We may consider the result of this sort of coherent summation as a *selective spectrum* of the power that is continuously distributed along the direction of θ_r. The denomination proposed is meant to recall one property of the ST: time averaging the local components collapses the TF representation into the global Fourier spectrum. The periodicity of the spectral components in $S_{\theta_r}[p']$ is the baseline to measure the repetition period T_0 of a sawtooth waveform made by chirps modulated with rate

$$\frac{B_{\text{eq}}}{T_0} = \tan(\pi - \theta_r)\,\frac{f_s^2}{N}. \tag{3.81}$$

The basic estimation of the periodicity consists of finding the amplitude peaks in the DFT. The spectrum can be paired with a detection threshold at constant false-alarm rate according to [159]. However, the temporal spacing between consecutive frequency bins exponentially grows with the period as a power of 2. Hence, the longer periodicities suffer from a coarse estimation resolution, while the shorter periods are precisely estimated, up to the higher bound defined by the Nyquist frequency. Another issue is the spectral leakage, namely the dispersion of the frequencies that are not integer multiples of any of the bins. As opposed to the DFT, the ACF can equally refine the estimation accuracy for both short and long periodicities, because it has constant temporal resolution. Nevertheless, this second function suffers from an excessive harmonic ambiguity: there exist peaks at the harmonics and the inter-modulation products of the actual periodicities. These spurious correlations likely complicate the choice of a significance threshold for detection, thus leading to many false alarms. More details about these methods are explained in [160], [161], and the references therein, for a variety of applications. Although DFT and ACF cannot separately provide reliable estimates of the periods, there are some approaches to combine them in a more consistent time representation in order to

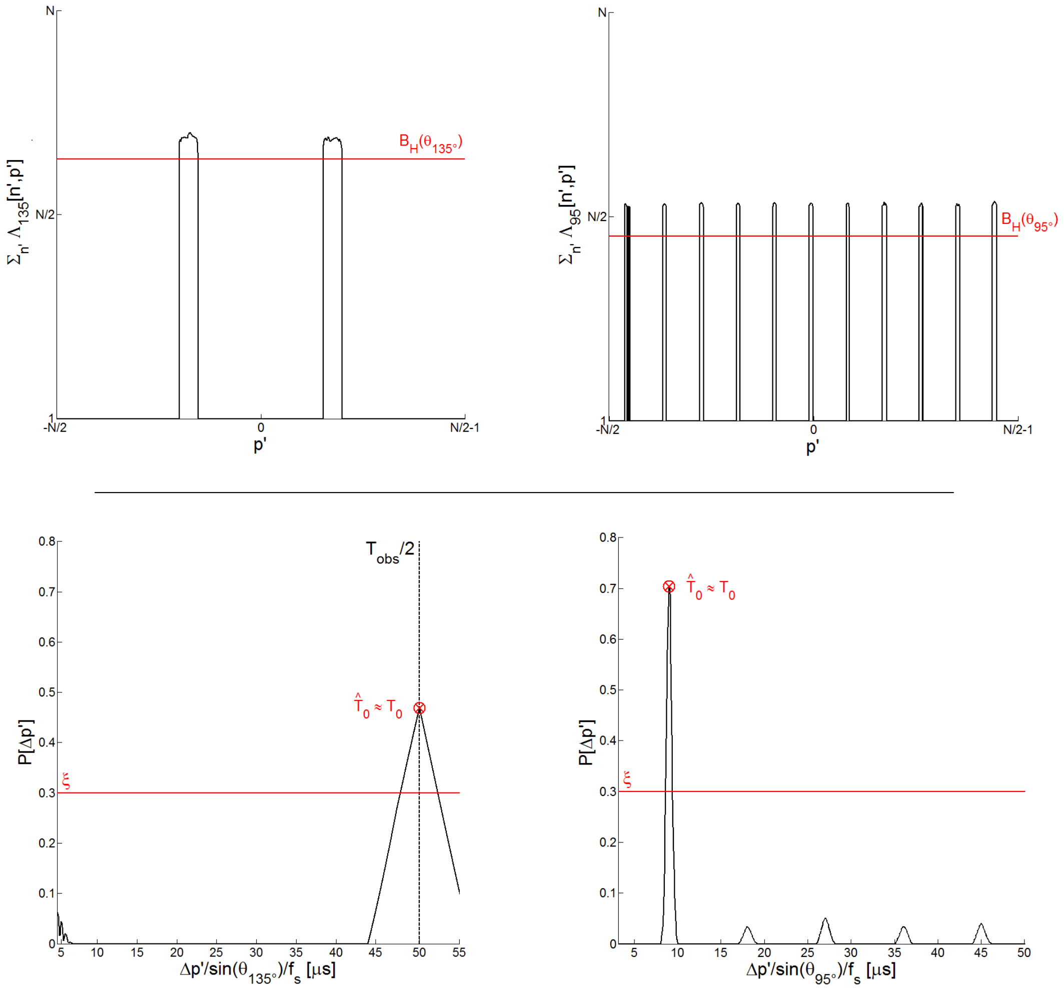

FIGURE 3.19: Outcomes of the chirp recognition and periodicity estimation applied to the energy mask in Fig. 3.18 for the two rotation angles $\theta_r = \{95°, 135\ °\}$ that roughly coincide with slopes of the two jammer slopes.

enhance both the periodicity accuracy and the robustness to both false alarms and missed detections. Recently, a candidate with these potential capabilities has been proposed in [162] for the classification of audio items based on the rhythm. Here, a fine-grained periodicity estimation is obtained as the maximum of a normalized and real-valued function, which merges the amplitude of the DFT and the ACF by interpolating and concatenating their time and frequency bins over a *hybrid lag-frequency axis* with index $\Delta p'$. The outcome of the interpolation is a periodicity function denoted by $P[\Delta p']$. For instance, applying this method to retrieve repetition periods highlighted in the rotated binary masks of Fig. 3.18 returns the estimates in Fig. 3.19. Besides the index corresponding to the maximum indicated in the pictures, all the bins with values above a certain significance threshold ξ are forwarded to the next stage of the algorithm. Since the periodicity function is normalized, the threshold can be easily designed.

Periodicity aggregation By repeating the previous two stages throughout the loop for all the angles of interest $\theta_r = \{95°, ..., 135°\}$, a two-dimensional periodicity

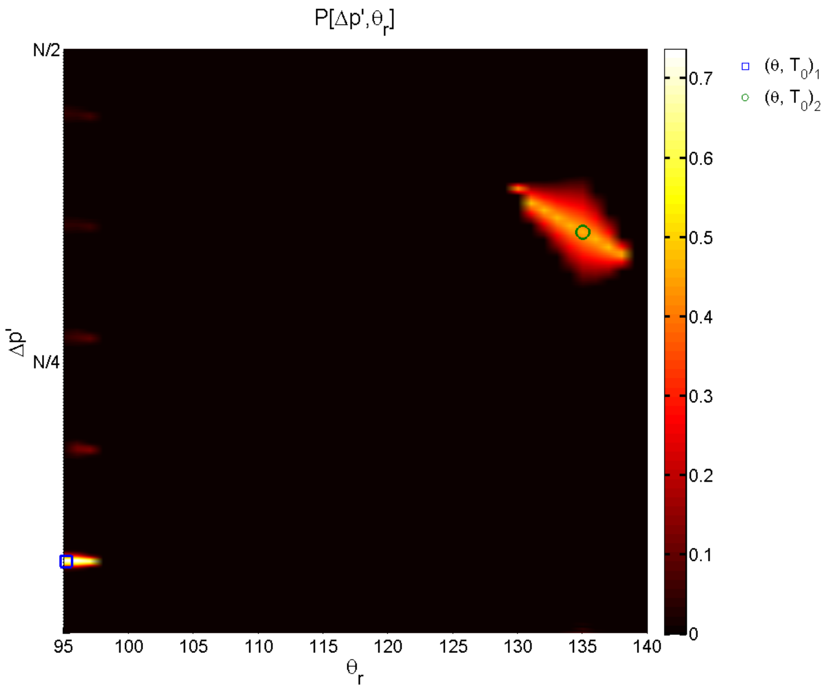

FIGURE 3.20: Periodicity map of the mask in Fig. 3.18 where the two jammers are characterized by $(\theta, T_0)_1 = (95°, 10\ \mu s)$ and $(\theta, T_0)_2 = (135°, 50\ \mu s)$.

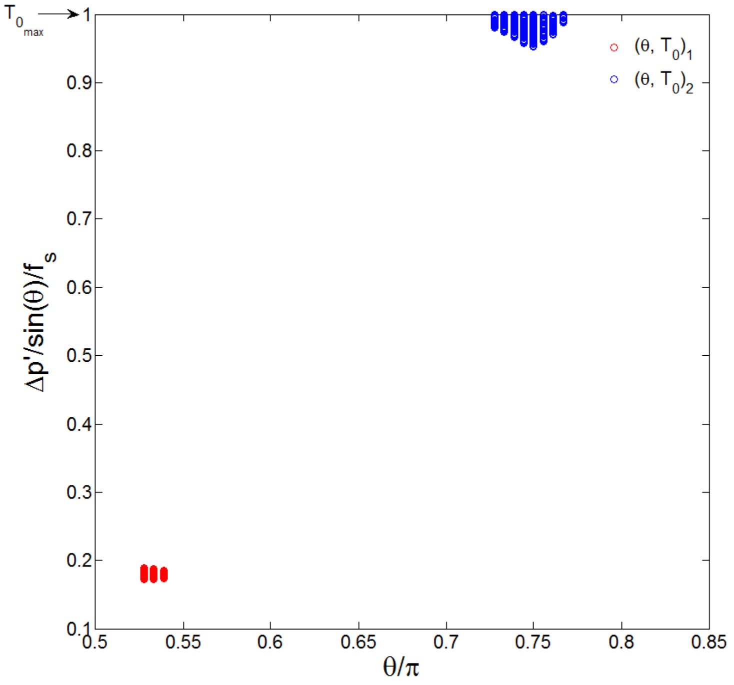

FIGURE 3.21: Clusters identified by the DBSCAN from Fig. 3.20.

map is obtained and denoted with $P[\Delta p', \theta_r]$, such as that in Fig. 3.20. As anticipated, the values in this map are normalized, thus setting a suitable threshold

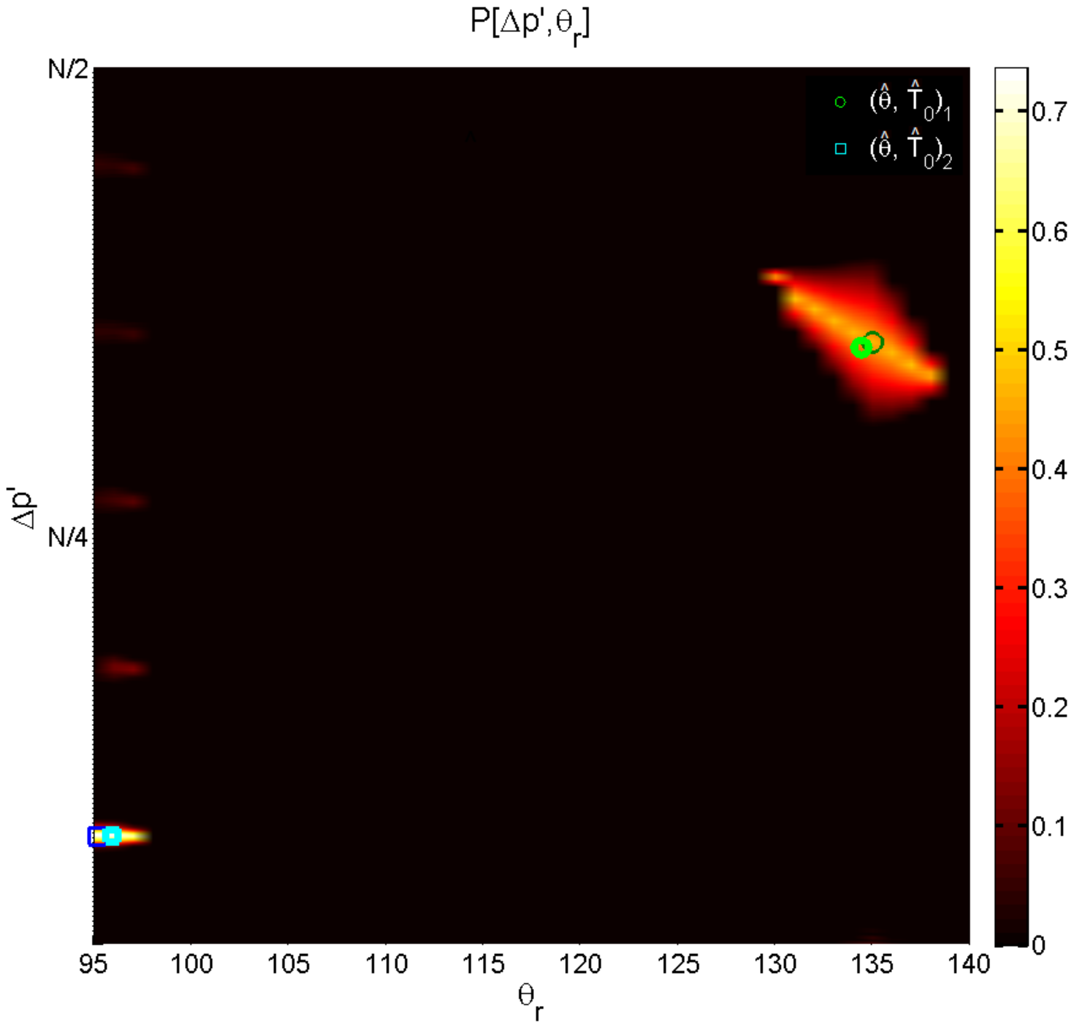

FIGURE 3.22: Pairs of jammer characteristics $(\hat{\theta}, \hat{T}_0)_1$ = (95.39°, 9.02 µs) and $(\hat{\theta}, \hat{T}_0)_2$ = (134.42°, 49.11 µs) estimated as the centroids of the clusters and projected onto the periodicity map of Fig. 3.20.

ξ is a matter of design. Generally, the more numerous are the waveforms contained in the snapshot and the closer are their characteristics, then the lower is the overall significance of the periodicities. Therefore, under a pessimistic circumstance, setting a loose threshold is preferable in order to avoid unrecoverable missed detections. For example, in Fig. 3.19 we set $\xi = 0.3$. Obliviously, this choice means that false alarms are more likely to occur, but they can be rejected by aggregating the periodicities estimated, as we do in the present third and last stage of the algorithm. The domain where we run the aggregation of the estimates is defined by normalizing the axes of the periodicity map. The significant periodicities are converted into the equivalent intervals of time between consecutive chirp repetitions through an angular factor (i.e., $\sin(\theta)$) and divided by the maximum period visible, namely $T_{0_{\max}}$. The chirp angles are normalized on 180°. In this new reference system, we process the map through a well-known routine that was named density-based spatial clustering of applications with noise (DBSCAN) in [163]. This tool efficiently aggregates density-connected estimates that are likely to be associated to the same jammer, while discarding the isolated ones that are probably false alarms. It requires two inputs: the maximum radius ϵ_{D} of two neighboring points to establish their direct connection and the minimum number of points N_{D} to form a cluster of sufficient density. Fig. 3.21 shows an example of clustering. The output is an estimate $\hat{N}_{\psi}$ of the number N_{ψ} of simultaneous jammers. Furthermore, the centroids of the clusters can be used to jointly estimate pairs $(\hat{\theta}, \hat{T}_0)_i$ of chirp rate and repetition period with $i = 1, ..., \hat{N}_{\psi}$. For the sake of clarity, Fig. 3.22 updates the map in Fig. 3.20 with the projection of the jammer characteristics extrapolated from Fig. 3.21. Given these estimates and thanks to the sawtooth nature of the jamming waveforms, one can also get an insight into the respective bandwidths:

$$\hat{B}_{\mathrm{RF}_i} = \hat{T}_{0_i} \tan(\pi - \hat{\theta}_i) \frac{f_s^2}{N} \tag{3.82}$$

even though the interference power spectrum might be well beyond the cutoff frequency of the front-end filter.

The three stages of the algorithm are summarized by the block diagram in Fig. 3.23, together with their inputs and outputs. The idea of resorting to TF analysis with the ST is promising, but this strategy also inherits a limitation: the characteristics of different jammers are supposed to be separated in terms of chirp slope. Two sawtooth waveforms with the same rate may be mistaken for one with a repetition period resulting from the combination of the two. The extent of the sufficient separation is inversely proportional to the TF resolution granted by the sampling rate and the observation time as well as proportional to the chirp slopes. Indeed, since the energy of the slower jammer is more spread in the TF domain, so are the periods associated to it in the periodicity map. This fact is noticeable in Figs. 3.19-3.22. Anyway, even with ideally infinite resolution, there must be *one and only one* jammer per possible chirp angle. This assumption is necessary to correctly estimate the periodicities on the map and to prevent DBSCAN from erroneously merging the periods estimated into a unique cluster. The impact of this limitation and possible countermeasures to it are discussed in the next subsection.

3.3.4 Numerical Results

From Fig. 3.15 we have derived the lower and higher bounds on the chirp angles that can be accurately estimated by analyzing the TF representation. Let us consider now a receiver observing a snapshot of 100 µs over a bandwidth of 20 MHz, which is that available in popular SDRs. In here, the range of values of θ covers just a portion of the characteristics featured by commercial in-car jammers and listed in Tab. 3.2. Particularly, it catches the average repetition periods and bandwidth (i.e., $T_0 = 10\,\mu\mathrm{s}$ and $B_{\mathrm{RF}} = 15\,\mathrm{MHz}$), but faster modulations lie under the lower bound of $\theta = 95$ and are thus poorly represented by the average-optimized ST. Likewise, the periodicity of any sawtooth waveform exceeding the higher bound of $\theta = 135$ is not visible. These shortcomings have opposite relevance. On the one side, according to the scientific literature, jammers slower than 50 µs barely exist on the market, because their quasi-stationary spectra are narrowband over long intervals of time. Hence, their interference is less effective and easier to filter at the receiver. On the other side, fast jammers should be addressed by shrinking the observation time, so that their chirp angles are steered until they enter the range of accepted values, namely $\theta \in [95°, 135°]$. Theoretically, the same can be accomplished by increasing

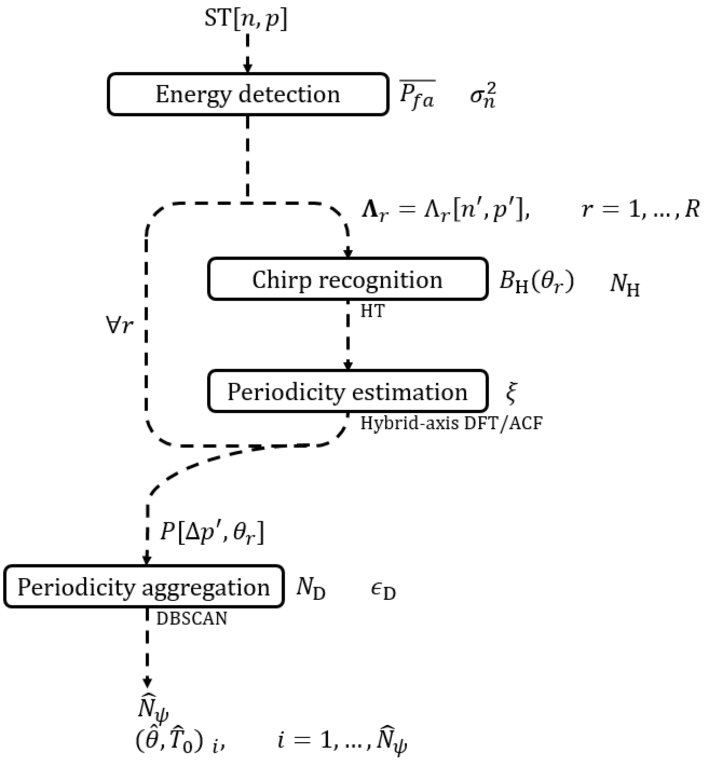

FIGURE 3.23: Block diagram of the algorithm for characterizing multiple in-car jammers by analyzing trains of linear chirps with the ST.

the sampling frequency, which is, however, fixed and unique for most of the implementations in practice. For instance, we depict in Fig. 3.24 a map of the "zones" that circumscribe the possible combinations of jamming characteristics, which we can estimate with a given equipment and consequent TF resolution. Each zone varies both f_s and T_obs, and so the complexity of the ST for analysis. The extension of the subspaces of sawtooth waveforms that can be searched by the algorithm changes accordingly. The aim in Fig. 3.24 is to cover the whole space with negligible overlaps between adjacent zones. More importantly, this example suggests that the algorithm is potentially able to tackle any incoming modulation, if we can sweep the rate of the ADC and/or the snapshot duration adopted by the receiver.

The analytic study of the estimation accuracy for the chirp angle and the repetition period is not viable, due to the large number of variables. The validation of the characterization algorithm is made by evaluating the performance in a few key situations with an extensive simulation campaign. The results in the presence of a single jammer are categorized into three zones that differ in the receiver configuration and that are numbered by (1), (2), and (3), as shown in Fig. 3.24. The arrays of simulations probe samples of waveforms that are indicated by the markers in the figure above and share the parameters listed in Tab. 3.3. Eventually, the estimation accuracy is tested with two jammers for the receiver configuration (1) only. The corresponding zone indeed occupies most of the search space of the jamming characteristics. Moreover, in parallel to the characterization, the long observation interval turns out ot be useful to estimate the attacker locations from precise TDOA measurements, as we illustrate in the last section of the chapter. These results give a deep insight into both the capabilities and the limitations of the algorithm presented.

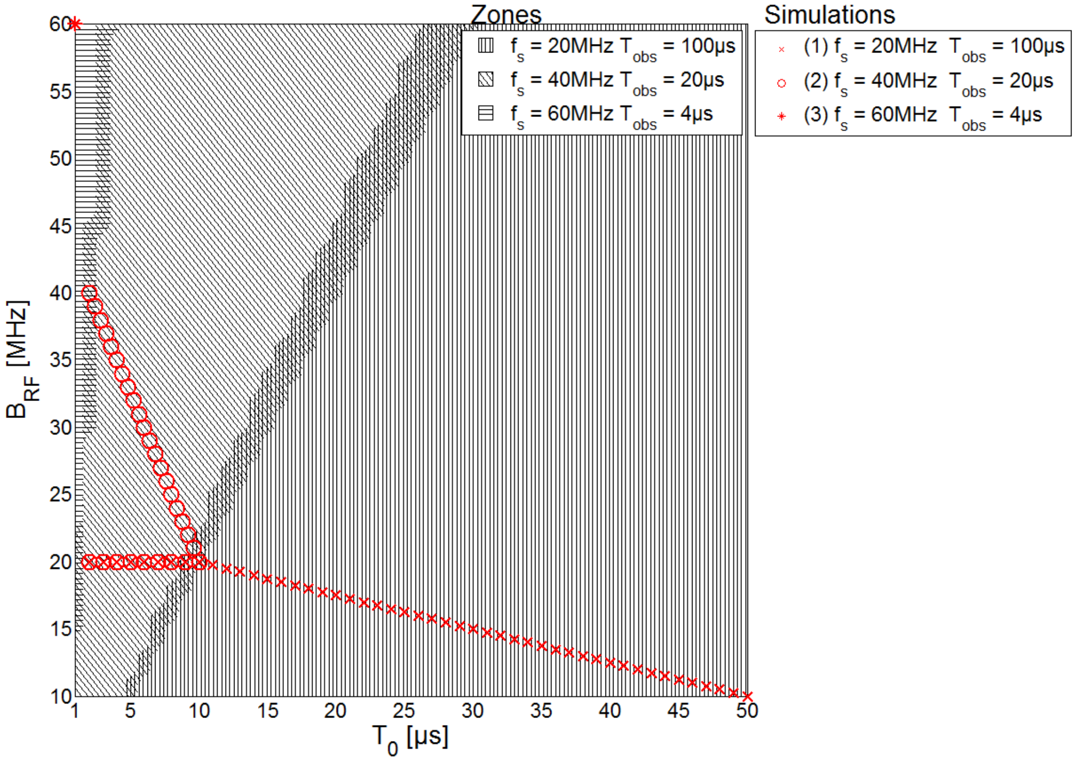

FIGURE 3.24: Combinations of repetition periods and bandwidths of the sawtooth that can be examined depending on the sampling rate and the observation time underlying the ST.

(1) Long Observation Time

We may obtain the accuracy of the estimation by assessing the estimation bias over a number of Monte Carlo realizations. The dashed lines around the error curves specify the 95th percentile of the results. The number of jammers estimated is steadily within the 10% of error if the JNR is greater than or equal to 9 dB, as shown in 3.25. The bias is contained also with lower power levels (e.g., 6 dB), when the jammer is not too fast (i.e., $\theta \geq 100\,°$). The performance of the estimates for the repetition period and chirp angles are remarkable, with average errors on the order of few tens of nanoseconds and one degree, respectively. In both, a sudden degradation is clear near the highest acceptable chirp angle, because the energy of the slower modulations is spread in the TF domain, notably with high JNR. Fig. 3.26 confirms the validity of the criteria behind the bounds on the chirp slopes: as expected, the algorithm does not work properly for jamming waveforms outside the zone of operation of the receiver used.

TABLE 3.3: Common parameters among the simulations of multiple jammer characterization.

Description	Symbol	Value
Single-sided noise power spectral density	N_0	-194 dBW/Hz
Front-end equivalent noise passband bandwidth	B_{eq}	0.5 f_s
In-band JNR level	ρ	$\{3, 6, 9, 12, 18, 24, 30\}$ dB
Energy detection false-alarm probability	$\overline{P_{\text{fa}}}$	N^{-2}
Maximum number of segments	N_{H}	$N/10$
Periodicity normalized significance threshold	ξ	0.3
Minimum cluster size	N_{D}	5
Maximum intra-cluster connection radius	ϵ_{D}	0.02

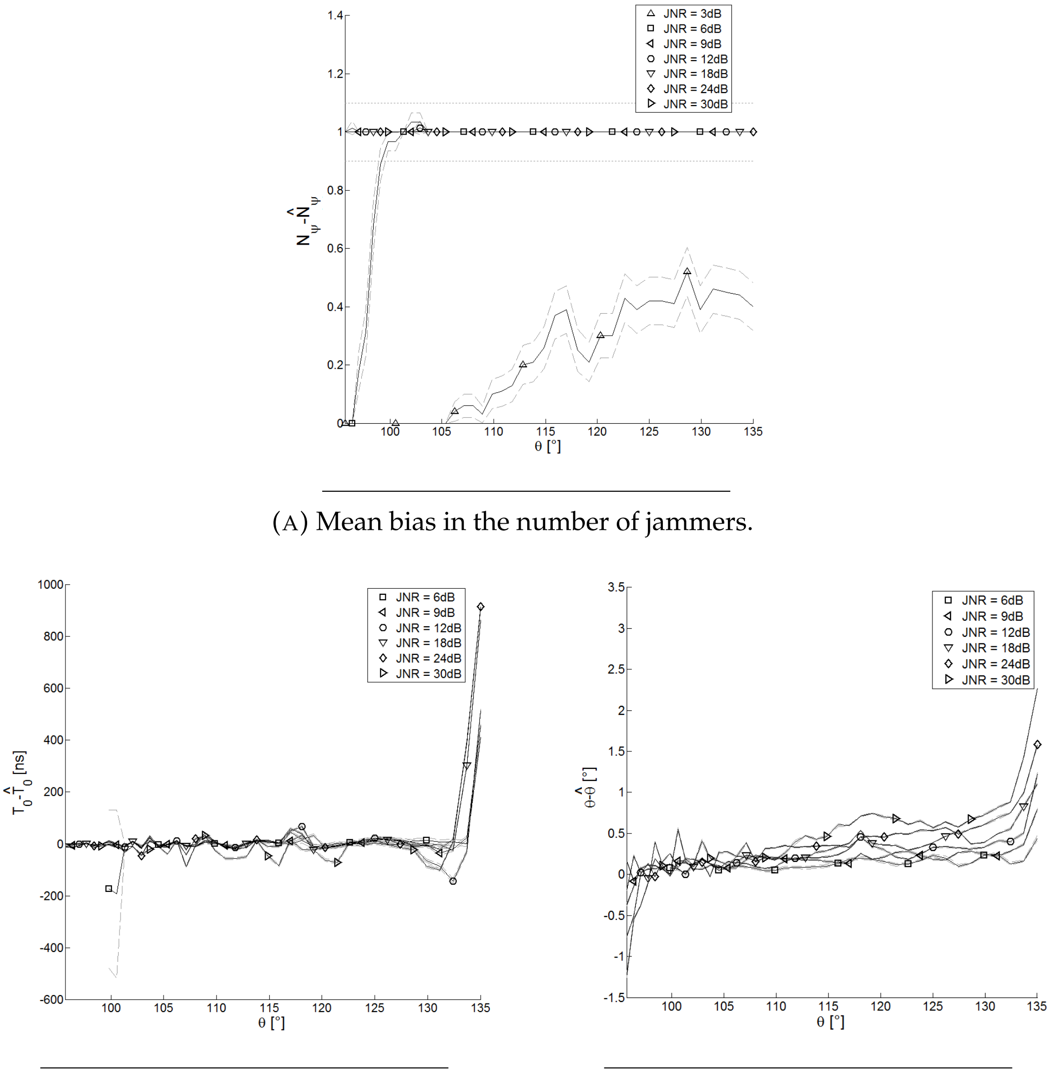

(A) Mean bias in the number of jammers.

(B) Mean bias in the repetition period.

(C) Mean bias in the chirp angle.

FIGURE 3.25: Estimation accuracy inside the zone for the receiver with $f_s = 20$ MHz and $T_0 = 100$ µs.

(2) Short Observation Time

The performance remain still satisfactory, exhibiting just faint signs of deterioration due to the lower TF resolution. The behavior with respect to the chirp angle is the

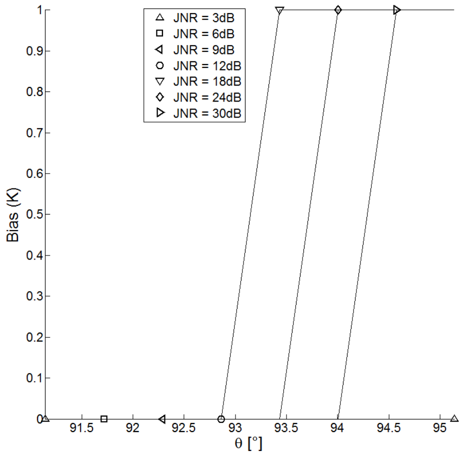

FIGURE 3.26: Mean bias in the number of jammers outside the zone appropriate for the received, but inside the adjacent one that would work with $f_s = 40$ MHz and $T_{\text{obs}} = 20$ μs.

same previously commented with the longer observation time. Fig. 3.27 explores the space with an interference bandwidth constantly equal to 20 MHz. We may compare the outcome to the same evaluation carried out with the wrong receiver configuration in 3.26. Contrary to this first array of simulations, a second one is executed by varying both the time and frequency characteristics of the jammer. The results in Fig. 3.28 are in line with the others and with the expectations. The estimates of the repetition period become more accurate, while the ones of the chirp angle are less.

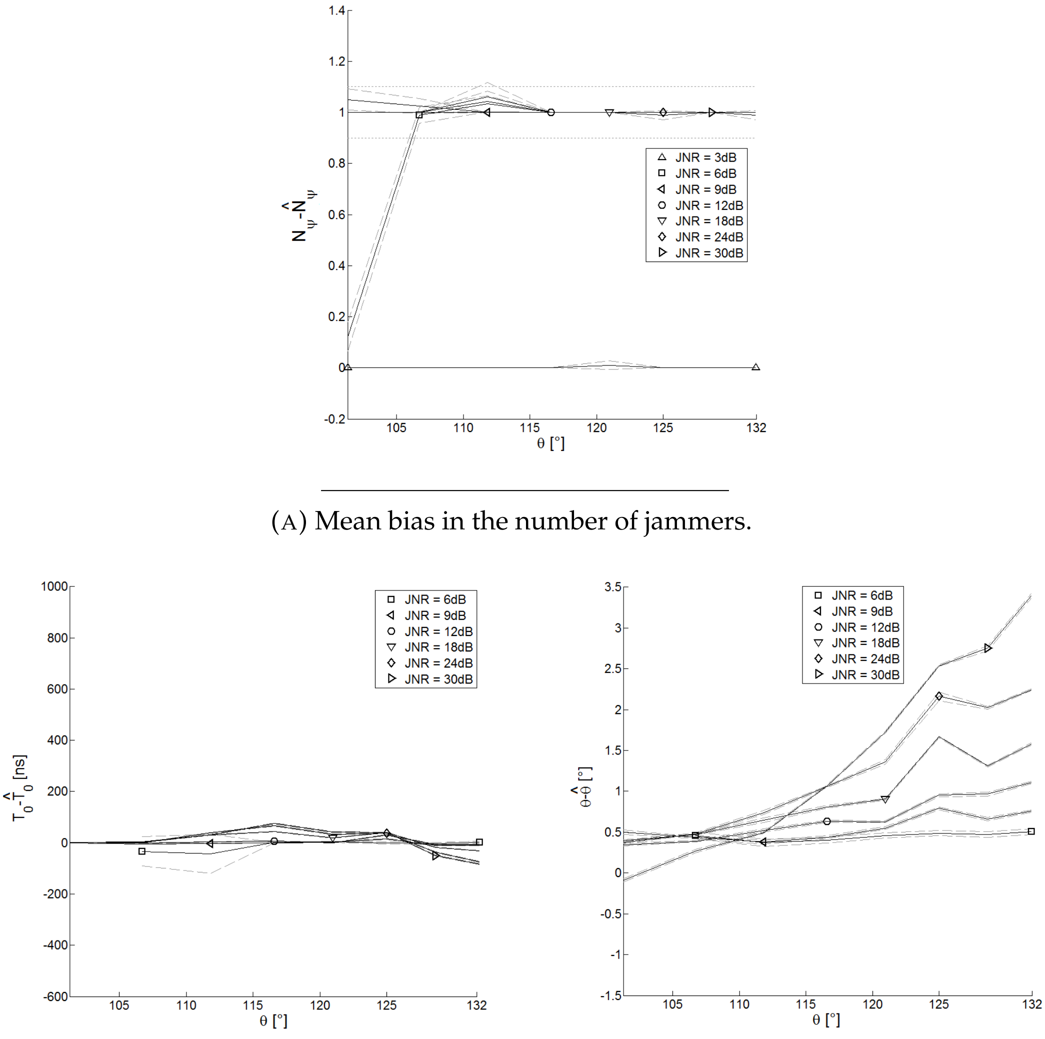

(A) Mean bias in the number of jammers.

(B) Mean bias in the repetition period.

(C) Mean bias in the chirp angle.

FIGURE 3.27: Performance for the first array of simulations run inside the zone for the receiver with $f_s = 40$ MHz and $T_{\text{obs}} = 20$ µs.

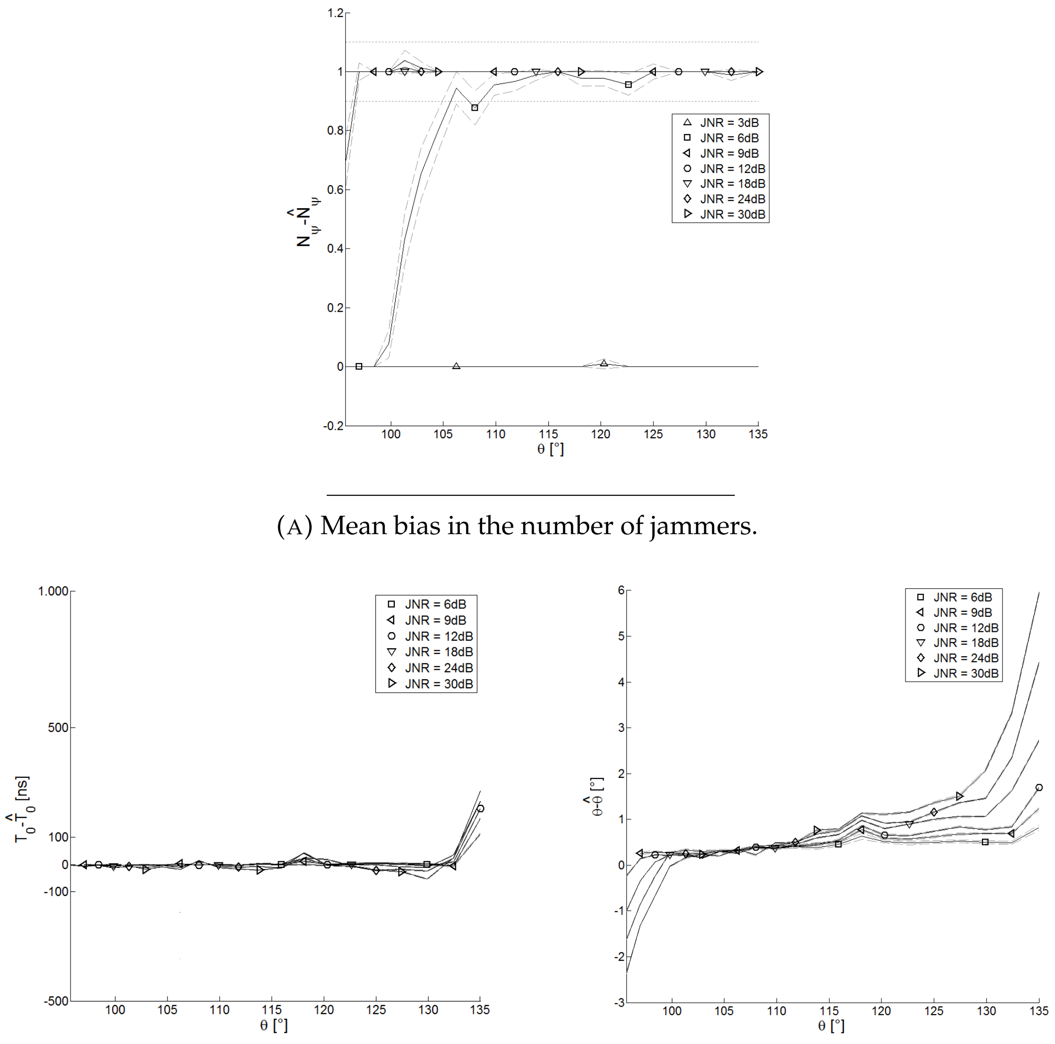

(A) Mean bias in the number of jammers.

(B) Mean bias in the repetition period.

(C) Mean bias in the chirp angle.

FIGURE 3.28: Performance for the second array of simulations run inside the zone for the receiver with $f_s = 40$ MHz and $T_{\text{obs}} = 20$ μs.

(3) Wideband Receiver

The last simulation with a single jammer tests the worst-case scenario where the sawtooth waveform is transmitted with $B_{\text{RF}} = 60$ MHz and $T_0 = 1$ μs, thus with the fasted modulation known. Although the TF resolution is way lower than before, the accuracy is still consistent with precedent results. As far as the angle of the chirp slope is concerned, the strong dependency on the JNR is due to the coarse granularity given by the sampling frequency and the observation time. When the interference power is weak, the performance are very sensitive to the few TF components detected. Conversely, if the jammer waveform is powerful at the receiver end, the detection of many components easily saturate the mask due to the poor resolution.

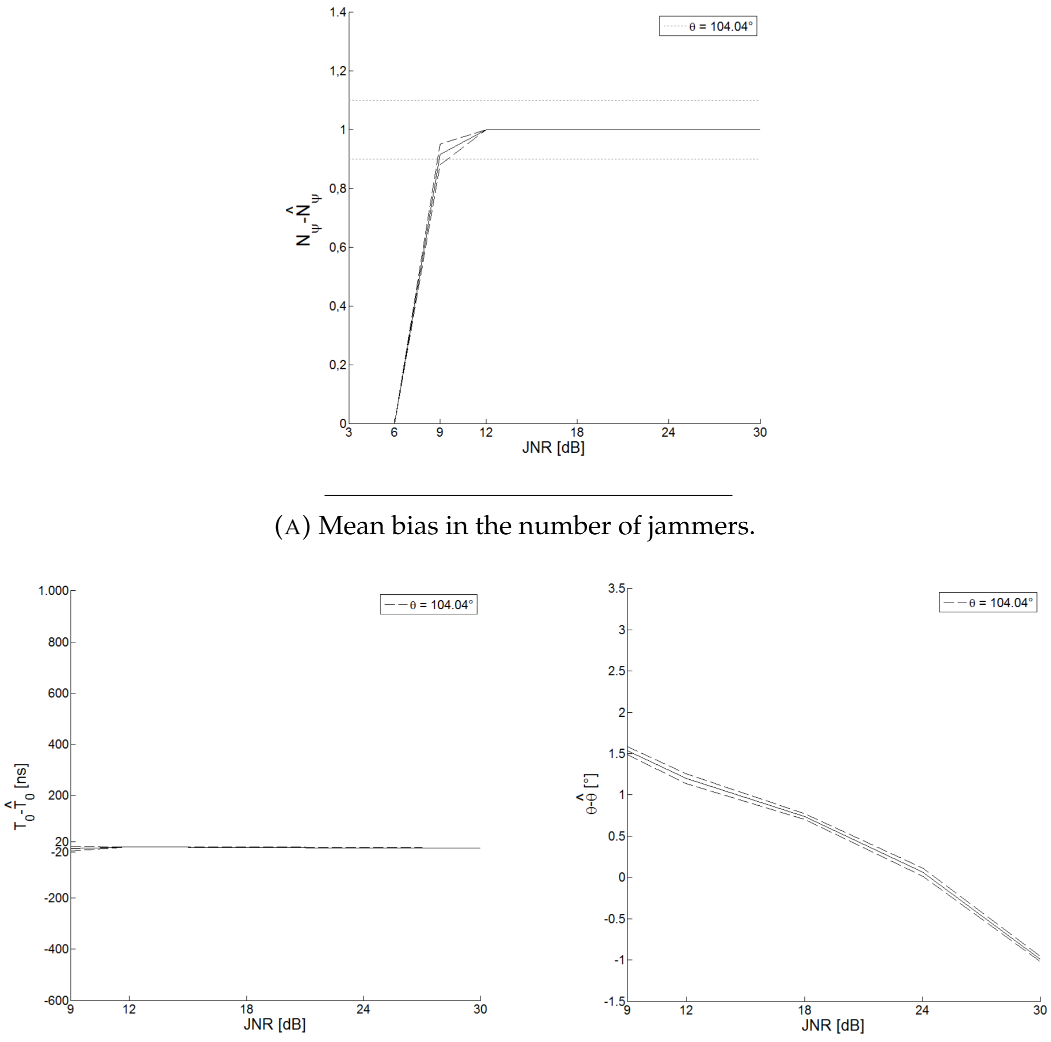

(A) Mean bias in the number of jammers.

(B) Mean bias in the repetition period.

(C) Mean bias in the chirp angle.

FIGURE 3.29: Characterization of the fastest jammer inside the zone for the receiver with $f_s = 60$ MHz and $T_{\text{obs}} = 4$ µs.

(1) Two Jammers

Since the performance for a single jammer is well surveyed, we may move on to a more interesting and realistic scenario: the reception of two jamming waveforms. The jammer characteristics are swept in the range of 10-20 MHz in frequency and between 10 µs and 50 µs in time, along the array of simulations marked in Fig. 3.24 for the receiver configuration (1). In view of the precedent results, we can restraint the simulation campaign to JNRs greater or equal to 9 dB in order to guarantee an accurate characterization of the individual sawtooth waveforms. The goal is to investigate how the algorithm is affected by the separation between the characteristics of the two sawtooth waveforms, depending on the respective received power levels. For this purpose, it is sufficient to test three combinations of JNR levels:

- $\rho_1 = \rho_2 = 9$ dB;
- $\rho_1 = 9$ dB, $\rho_2 = 30$ dB;

- $\rho_1 = \rho_2 = 30$ dB.

In the first scenario, both the jammers are sufficiently powerful to allow for a TF representation that is robust against noise, but not enough to cause the saturation of the mask computed on the energy in the ST. This is then a best-case scenario. The results in Fig. 3.30 highlight both the achievements and the setbacks of the algorithm. If the number of jammers is correctly guessed by the DBSCAN, the periodicity and the slope are precisely characterized. The performance in terms of RMSE for $\hat{T}_0$ and $\hat{\theta}$ are shown in Figs. 3.30b and 3.30c only where the mean error in $\hat{N}_\psi$ is lower than 10%. Remarkably, the accuracies are below the microsecond and one degree, on average. However, they depend on the correct separation of the two jammers, which is challenging when the chirps they transmit are modulated by a similar slope. Indeed, the bias in $\hat{N}_\psi$ statistically grows far beyond the 10%, if the sawtooth waveforms feature a differential chirp angle up to ±5°. Such a limitation is more evident for the other two cases in Figs. 3.31 and 3.32. Here, the situation is worse due to the higher strength of the jammers. In fact, since more energy is captured, then more components are detected and the mask is prone to saturation. As a consequence, the resolution of the TF representation searched for periodicities may be insufficient for the algorithm to distinguish among mixed jamming waveforms.

In order to enhance the effectiveness of the algorithm, it is necessary to introduce some form of *power adaptation* into the receiver, akin to the AGC loop. As a matter of example, let us consider the performance in Fig. 3.32. Intuitively, we can force the algorithm to achieve accuracies near those in Fig. 3.32 with a simple workaround: we lift the noise floor for the energy detection in the ST. More specifically, we may add about 21 dB to the noise power σ_n^2, which goes from 30 dB to 9 dB below the interference, so that the resultant binary mask solely includes the most significant TF components. Generally speaking, we may think of σ_n^2 as a "knob" for adjusting the sensitivity of the detection stage. By tuning this quantity, we can raise the threshold in Eq. 3.75 for all the frequencies at once in order to prevent saturation and aid the separation of the slopes if needed. This adaptation is not effective when there is a substantial gap in terms of received power between the two jamming waveforms. This is the case of Fig. 3.31. Nevertheless, we can turn this issue into an advantage by adapting the detection threshold to characterize the strongest jammer first, while the weaker one is hidden. Afterwards, the components of the train of chirps are canceled out and the threshold is restored to a lower value. Now, we can re-run the search routine to focus on the characterization of the weakest jammer. This strategy of *successive cancellations* is exploited in the next section to extrapolate the waveforms and compute separate CAFs.

3.3.5 Conclusions

The goal behind the previous numerical evaluation is to demonstrate the promise of using the typical periodicities of the sawtooth pattern that characterizes the frequency modulation of jamming waveforms in order to count and identify the respective sources. Particularly, we initially argue that the ST is the proper tool to carry out the TF analysis of unknown signals, since the receivers under attack do not know beforehand the characteristics of the jamming waveforms, and more specifically the bandwidths and the repetition periods. Therefore, through the ST, the algorithm proposed shifts the analysis from the delay-Doppler domain (i.e., the CAF) to a new domain in the TF plane, where separating the waveforms becomes possible. This accomplishment is obtained by discriminating the periodicities of sawtooths. What is crucial in here to understand is that different jammers are assumed to generate chirps with different slopes; this assumption is necessary unless other conditions on their received power levels take place. In Fig. 3.14, the chirp slope is defined as an angle that depends on the bandwidth and the repetition period used to modulate the sawtooth, and also on the TF resolution tied to the sampling rate of the receiver. Therefore, by properly tuning the sampling frequency and/or the observation time, the algorithm is able to distinguish jammers that transmit chirps with any realistic slope, as long as the TF resolution is sufficient. This fact is illustrated in Figs. 3.15 and 3.24. On the contrary, when the aforementioned assumption is relaxed, even in the new domain a pair of jammers might be hardly recognizable and mistaken for the same interference source, if their sawtooths have two angles separated by a few degrees (e.g., $\pm 5°$). The importance of such a limitation is assessed with a plethora of cases that consider two jammers and a reasonable TF resolution (i.e., that of a front end with 20 MSPS and 10-MHz bandwidth) that is within the reach of DoOs. The outcome of the simulations offers an insight into the performance achievable with any number of pairs of simultaneous jammers, which can be so characterized unless the TF resolution is too coarse or their respective chirp slopes are indistinguishable. Interestingly, as anticipated, a diversity in the energy captured from the two jammers might come helpful to overcome the lack of resolution necessary to distinguish them based on their inner periodicities. In brief, this possibility requires adjusting the sensitivity of energy detection in the ST, while successively analyzing and canceling sawtooth waveforms of decreasing power.

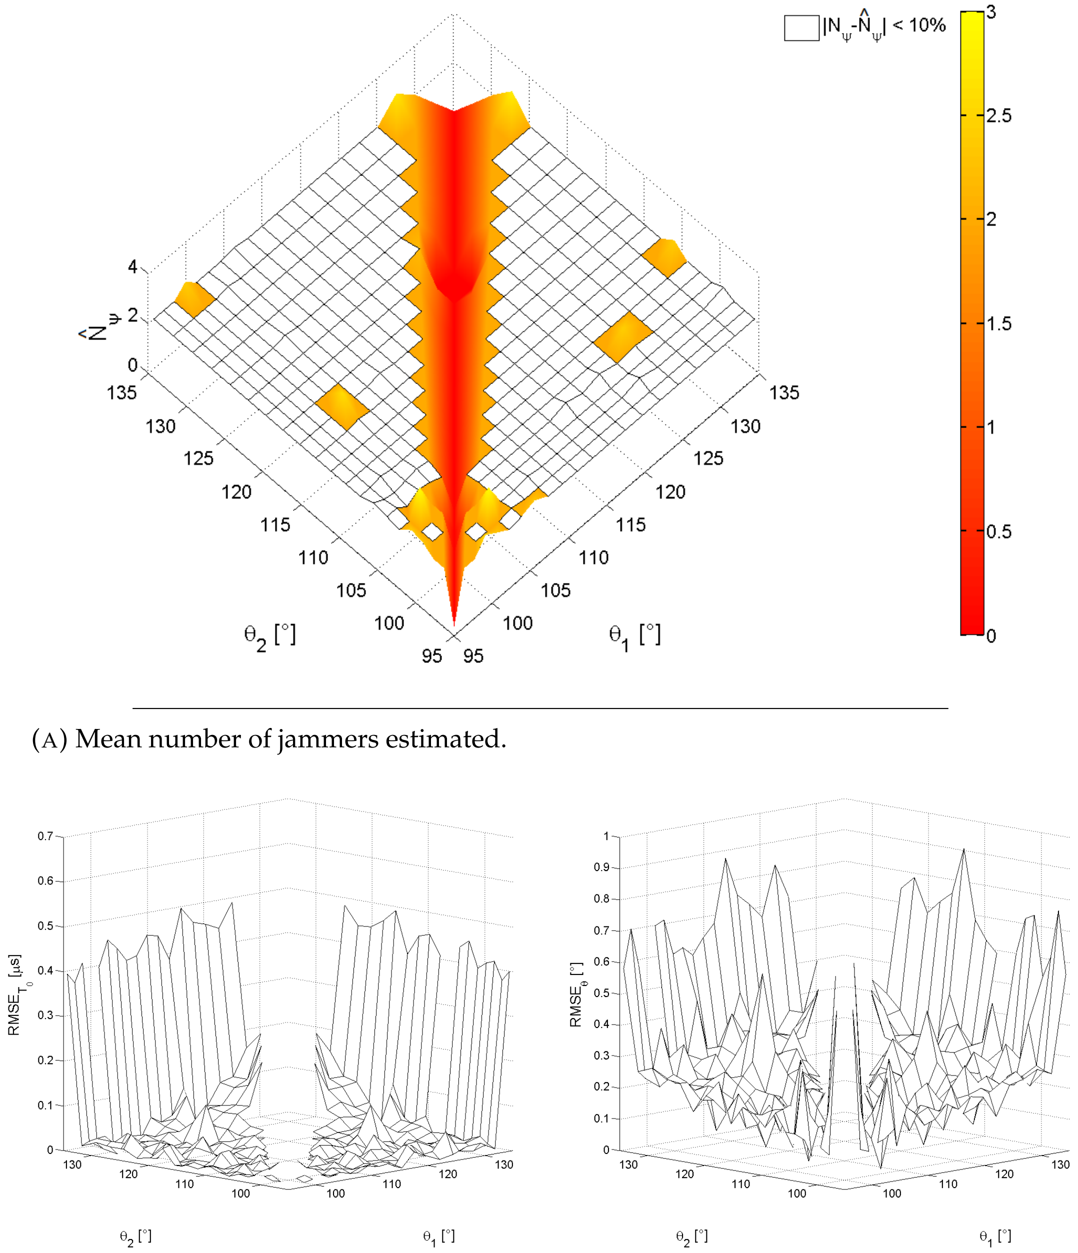

(A) Mean number of jammers estimated.

(B) Mean bias in the repetition period.

(C) Mean bias in the chirp angle.

FIGURE 3.30: Characterization of two jamming sawtooth waveforms having both JNR equal to 9 dB.

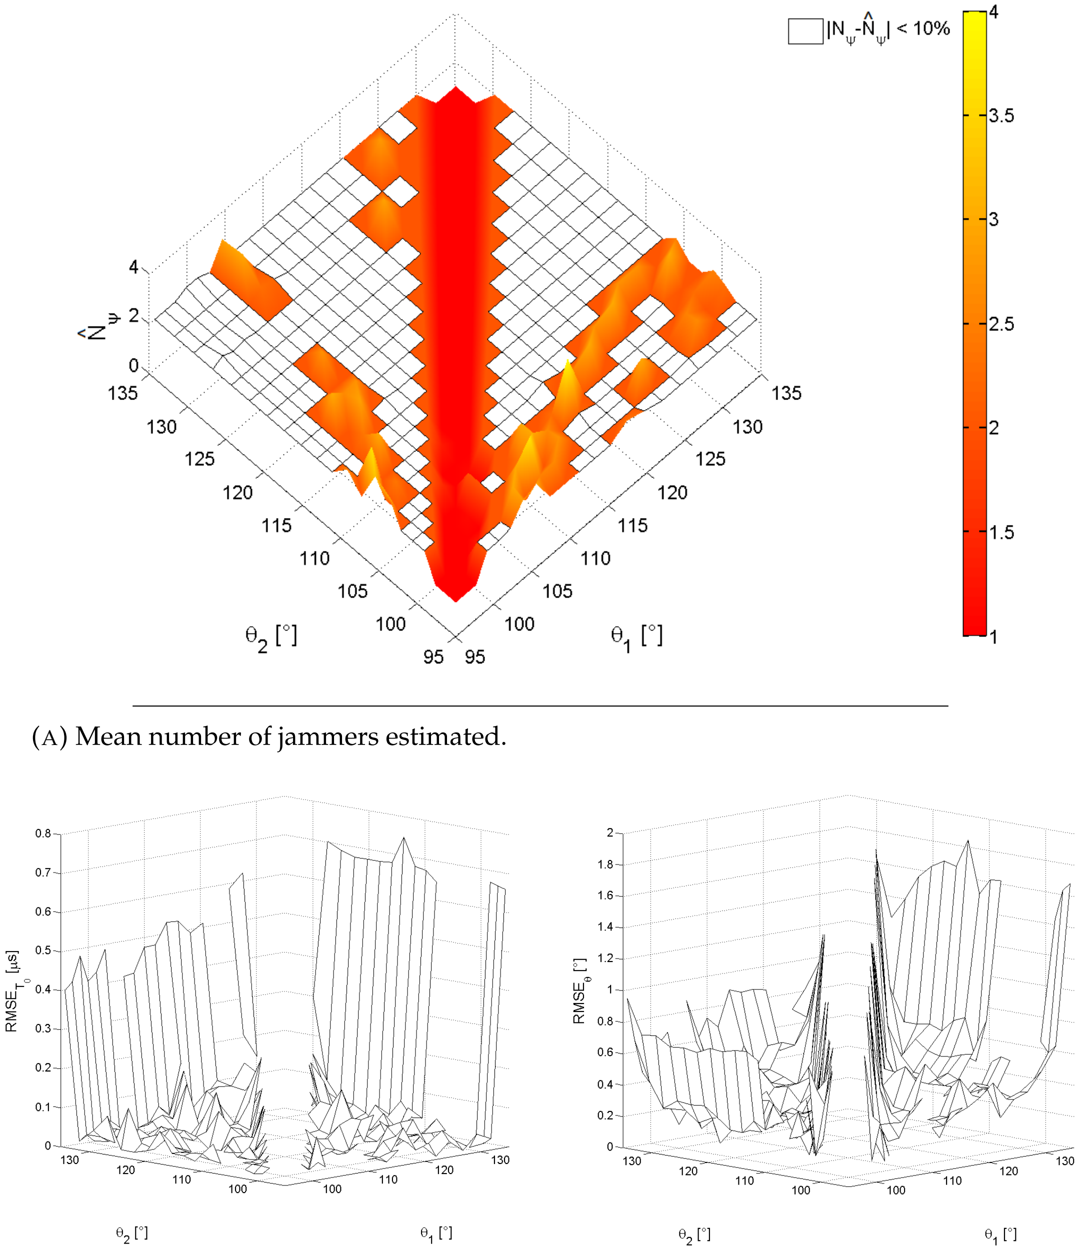

(A) Mean number of jammers estimated.

(B) Mean bias in the repetition period.

(C) Mean bias in the chirp angle.

FIGURE 3.31: Characterization of two jamming sawtooth waveforms with JNRs equal to 9 dB and 30 dB.

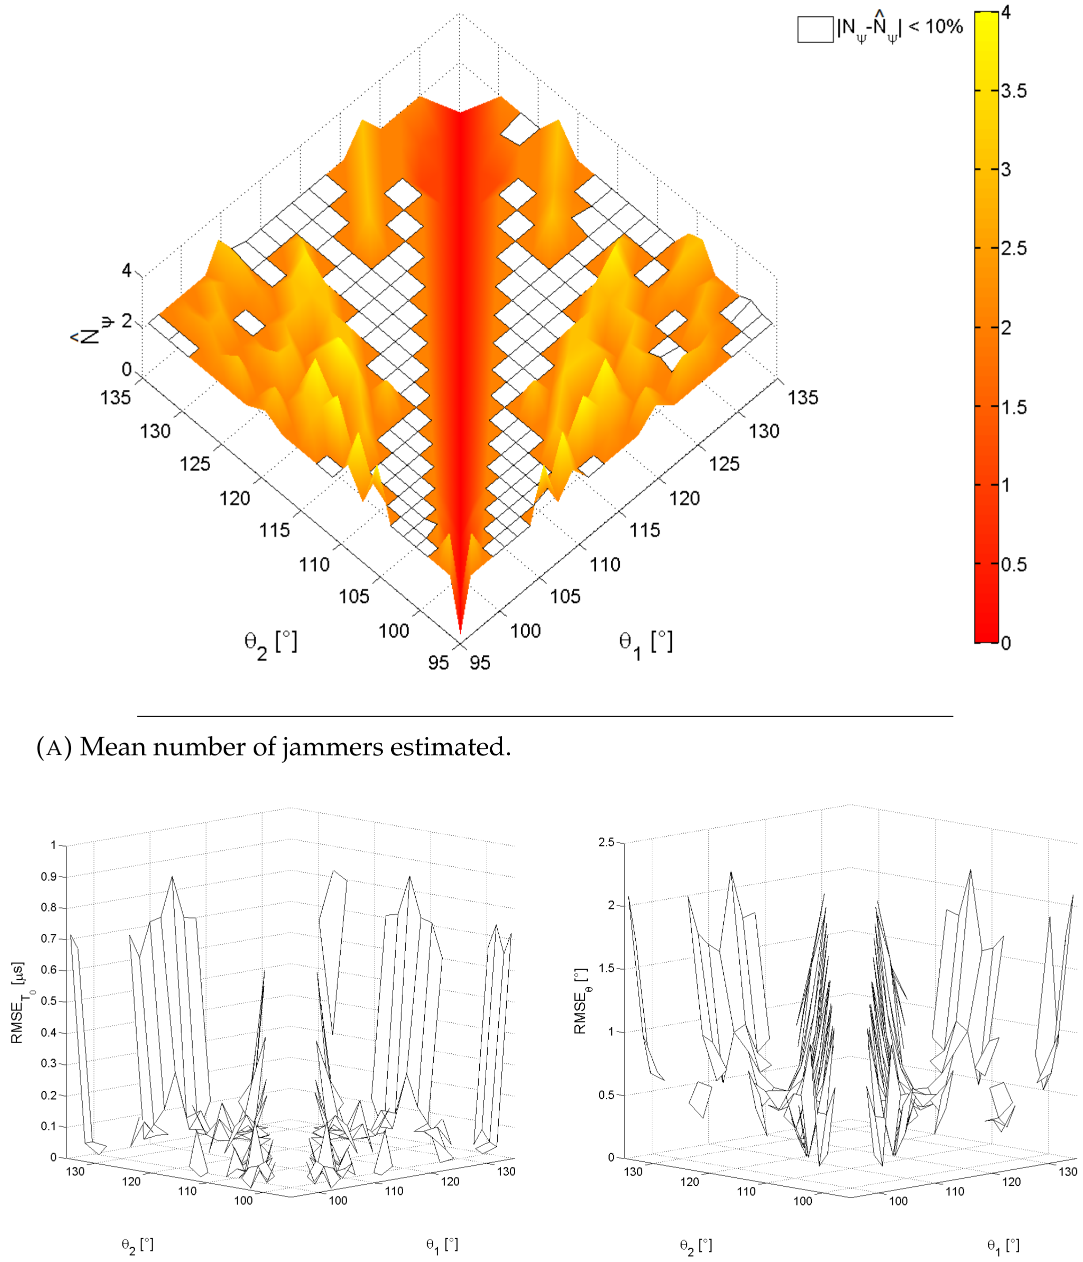

(A) Mean number of jammers estimated.

(B) Mean bias in the repetition period.

(C) Mean bias in the chirp angle.

FIGURE 3.32: Characterization of two jamming sawtooth waveforms having both JNR equal to 30 dB.

3.4 Snapshot Tracking of Multiple In-Car Jammers

The present section wraps up the methods previously presented in this chapter in order to finally track multiple jamming attacks at the same time. The goal is to enforce prompt actions against fast-moving in-car jammers, by passively tracking the positions and velocities of their respective origins within the area monitored by DoOs. The enabling idea at the base of such an application is to distinguish the jamming waveforms in the TF domain by means of their characteristic signatures. For this purpose, the characterization algorithm described in the previous section can do the job. As already explained, it leverages on the typical sawtooth pattern used to modulate interference: the incoming waveforms are expected to be a mixture of trains of linear chirps. This supposition is reasonable, because it is coherent with the signals generated by the vast majority of devices surveyed in [13, 14, 146]. The range of characteristic bandwidths and repetition periods are summarized in Tab. 3.2. The main assumption for this approach is the same necessary for characterizing the jammers: the chirp slopes should uniquely identify their respective sources, at least whenever their JNR levels are comparable. Under this circumstance, the discrimination of superimposed jamming waveforms based on their TF characteristics enables the extension of the method for the snapshot localization of a single source of interference to scenarios with a unknown number of jammers. In other words, the two-step estimation process in Fig 3.2 can separately operate on multiple jammers, by splitting the CAF into the waveforms separately characterized at the receiver end. This operation works in a *divide-and-conquer* fashion. The great advantage is that we can bypass the two main challenges behind this application: the joint estimation of concurrent TDOA/FDOA pairs and the association of the pairs estimated for separate receivers to the corresponding jammers. For the sake of clarity, these challenges are exemplified in Fig. 3.4.
The content of this section represents an original contribution to the state of the art and is being submitted for publication in [164].

3.4.1 Divide and Conquer

The approach proposed in the following allows us to address the presence of an unknown number N_ψ of jammers by splitting the problem into N_ψ sub-problems, the solution of which is more practical.

Disambiguation of Multiple Cross-Ambiguity Functions

When two or more jammers are simultaneously present in the area monitored, their CAFs interact constructively and destructively, depending on the relative distances between them and each pair of DoOs. In fact, as the jammers move, their peaks in the overall CAF can affect each other, thus causing false alarms as well as missed detections. The first problem here is reliably estimating the TDOA/FDOA pairs, given the fact that a ML is not a suitable solution. Secondly, a data association problem arises: the observables should be fed to the correct estimator, namely to that tracking the corresponding jammer state. To overcome both these challenges, we should gain insight into the jamming waveforms contained in every set of snapshots that is captured by a number M of separate DoOs. For this purpose, we may inspect the current snapshots through the ST optimized with Eq. 3.72 and then execute the characterization algorithm described in Fig. 3.23. Afterwards, at the time epoch k, the resultant set of periodicity maps $P[\Delta p', \theta_r]_{i_k}$ with $i = 1, ..., M$ can be averaged into

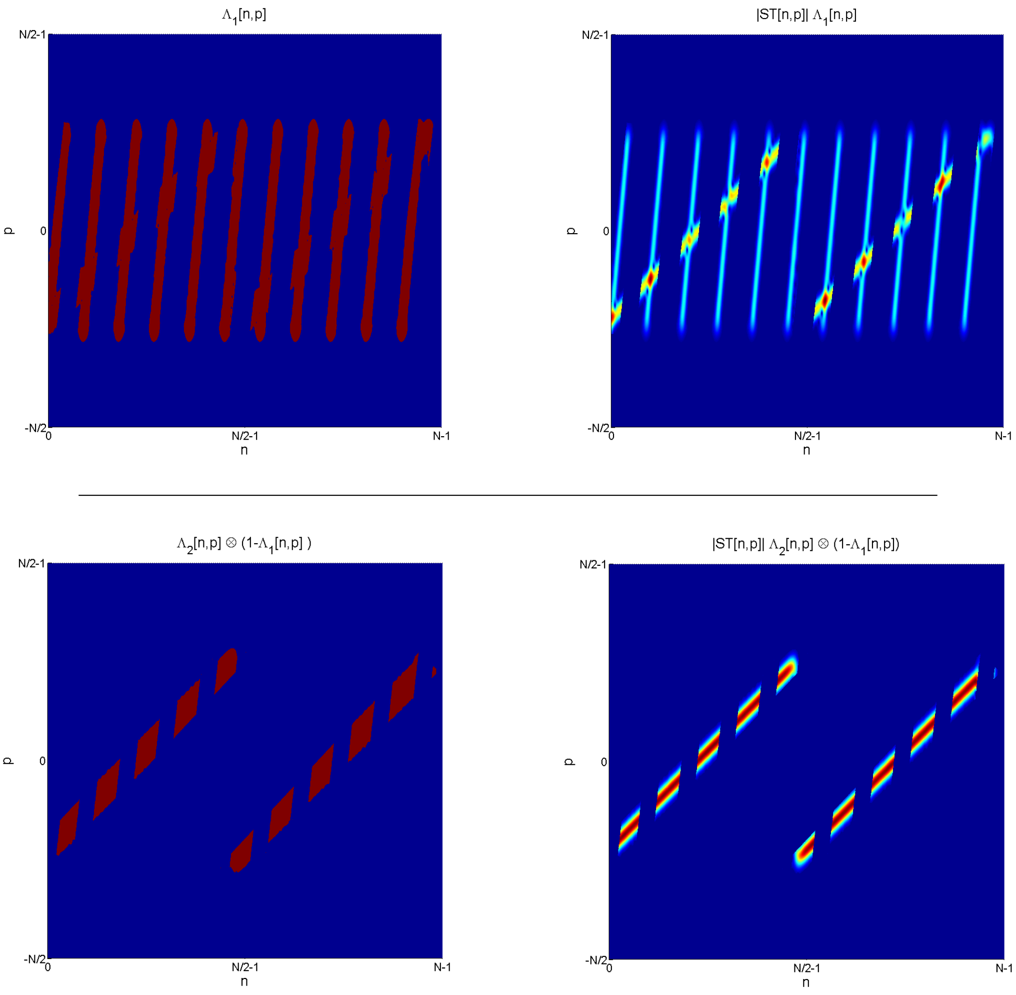

FIGURE 3.33: Separation the jamming sawtooth waveforms in Figs. 3.16b and 3.17 according to the estimated characteristics.

one (i.e., $\overline{P}[\Delta p', \theta_r]_k$), which is robust to missed detections and false alarms thanks to the spatial diversity of the receivers. By clustering the significant periodicities in $\overline{P}[\Delta p', \theta_r]_k$, we finally obtain a triad of estimates $(\hat{\theta}_j, \hat{T}_{0_j}, \hat{B}_{\mathrm{RF}_j})_k$ for the j-th jammer, with $j = 1, ..., \hat{N}_{\psi_k}$. The quality of these characteristics is crucial to disambiguate the individual CAFs.

Separate Reconstruction of Sawtooth Waveforms

The two-dimensional function $\overline{P}[\Delta p', \theta_r]_k$ mirrors the distinct periodicities recognizable in the TF domain. The combination of the periodicity maps overcomes the inner ambiguity of the delay-Doppler domain underlying multiple CAFs, when many jamming waveforms are received. We can exploit the knowledge gained from characterizing the jammers to extrapolate the individual jamming waveforms into the current snapshots, based on their TF characteristics. More specifically, we derive and apply a number $\hat{N}_{\psi_k}$ of binary masks $\mathbf{\Lambda}_{j_k}$ that act as TF filters, each of which lets only the components with significant energy aligned to the slope of $\hat{\theta}_j$ through periodic slots of length and width denoted by $B_{\mathrm{H}}(\hat{\theta}_{j_k})$ and σ_{IF}, respectively. The coordinates for centering the slots are inferred from the HT computed by the characterization

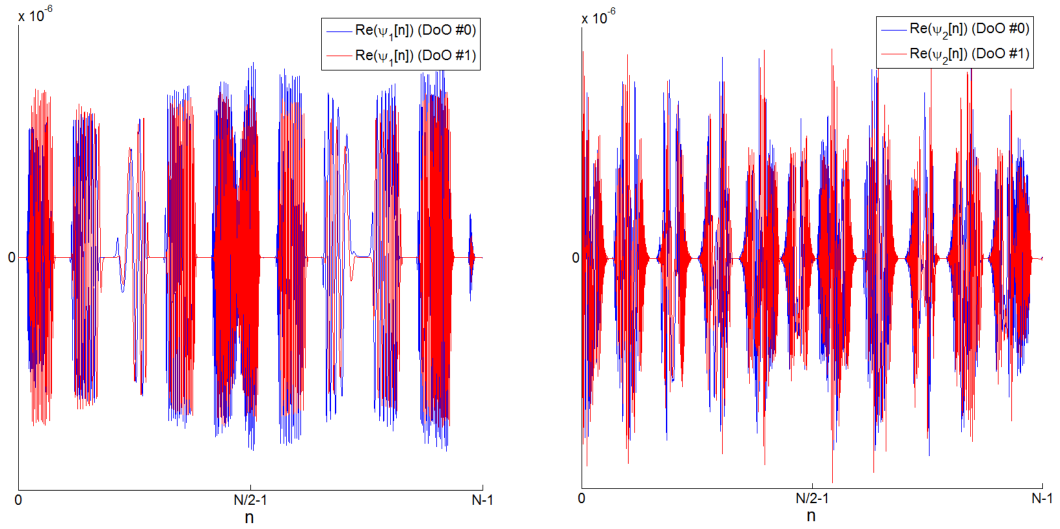

FIGURE 3.34: Jamming sawtooth waveforms reconstructed from the separation of the replicas, such as that in Fig. 3.33, which are contained in the snapshots received by two DoOs at the same time.

algorithm, as shown in Fig. 3.18. Their width should be proportional to the energy dispersion around the IF. As such, whilst arbitrary in practice, σ_{IF} would be theoretically adapted to every jammer as a function with respect to the estimates of chirp slope and JNR. The masks are used to separate the jamming waveforms by adopting a strategy of *successive cancellations* in descending order according to the received power. This procedure is necessary to filter out the mutual interactions among the received waveforms: the amplitude and the phase of certain frequency components can hide the concurrent contributions of multiple jammers. The cancellation is performed iteratively with $\hat{N}_{\psi_k}$ iterations. At first, the components associated to the highest JNR are simply extracted by multiplying the mask and the ST. From the second iteration on, the summation of the masks at precedent iterations for higher JNRs are canceled from the latest mask in order to remove the components evidently affected by the more powerful received interference that comes from other sources. After each cancellation, the mask results from binary additions and multiplications in the form of $\boldsymbol{\Lambda}_{j'_k} \otimes (1 - \boldsymbol{\Lambda}_{1_k} \oplus ... \oplus \boldsymbol{\Lambda}_{j'-1_k})$, given that the index j' sorts the jammers in order of increasing JNR (i.e., $\rho_{j'_k} \leq \rho_{j'_k-1}$) with $j' = 1, ..., \hat{N}_{\psi_k}$. The procedure described is clear from Fig. 3.33, where the TF components of two jammers received with comparable power levels are extracted through two iterations.

Once the sawtooths visible in the current set of snapshots are properly separated, the jamming waveforms captured by various DoOs can be reconstructed by inverting the STs filtered through successive cancellations. For the reasons discussed in detail in Chapter 2, the reconstruction is carried out with the TI shown in Fig. 2.6. Let us consider two distant and synchronous receivers, then the TF components extrapolated from their snapshots, such as the one in Fig. 3.33, are replicas of the same sawtooth waveform with different delays and Doppler shifts. They are plotted in Fig. 3.34 for instance. Obviously, after the extraction, the trains of linear chirps are irremediably distorted by both the TF filter and the unfiltered components related to the interactions with other waveforms at the reception. Nevertheless, their erroneous approximation is not important, because we are more interested in isolating the delays and the Doppler frequency shifts among the replicas reconstructed,

TABLE 3.4: Parameters for the case study of multiple in-car jammers.

Description	Symbol	Value
Carrier frequency	f_c	1575.42 MHz
Sampling frequency	f_s	20 MHz
Single-sided noise power spectral density	N_0	-194 dBW/Hz
Front-end equivalent noise passband bandwidth	B_{eq}	10 MHz
Jammer passband bandwidths	B_{RF}	$\{20, 15\}$ MHz
Jammer repetition periods	T_0	$\{9, 13\}$ µs
Jammer chirp angles	θ	$\{95.1, 99.8\}$ °
Jammer transmit powers	P_t	$\{0, \text{-}20\}$ dBW
Observation time	T_{obs}	0.1 ms
Normalized snapshot rate	R_s	1/5000
Snapshot period		0.5 s
Simulation time	T	4 s+T_{obs}
Number of observations	K	40001
Number of snapshots		9

rather than re-generating an exact copy of the jamming waveforms. Indeed, the ultimate goal is cross-correlating the replicas of the same sawtooth from two different DoOs in order to provide a CAF with a sufficiently clean main lobe. The replicas are associated to each other coherently with the characteristics estimated, namely by pairing the estimates $(\hat{\theta}_j, \hat{T}_{0_j})_{i_k}$ for $i = 1, ..., M$. Once their association is done, they can be used to compute separate CAFs, which are ideally *independent*, as if they were referred to single sources in different scenarios instead of multiple jammers in the same scenario. That is the basic concept behind the divide-and-conquer strategy. Given 4 DoOs and 2 jammers, this procedure leads to the parallel computation of a total of 6 CAFs per snapshot. Now, the ML criterion is supposed to provide a reliable estimation of the TDOA/FDOA pairs: the maximum of the main lobes of every CAF provides observables robust to the presence of multiple in-car jammers of diverse characteristics. Since the data association problem is solved, these estimates can be directly used as measurement into the EKF1 that tracks the corresponding origin. Therefore, given N_ψ jammers, there are an equal number of parallel trackers. In reality, an array of recursive estimators could be dynamically adapted to the jammers detected and characterized in the last sets of snapshots. In the following, the detection is given for granted and always asserted, because the interference is assumed to be transmitted with continuity.

3.4.2 Numerical Results

While an extensive campaign would be more convincing but less practical, we decide to evaluate the effectiveness of the divide-and-conquer strategy together with the jammer characterization algorithm for just one challenging case study. The rationale behind this choice is that the characterization capabilities of one single receiver have already been extensively studied in the precedent section, and the estimation of the locations of two jammers is straightforward when they are correctly characterized. In any case in order to get meaningful results through Monte Carlo simulations, the testing scenario is designed to stress both the localization accuracy and

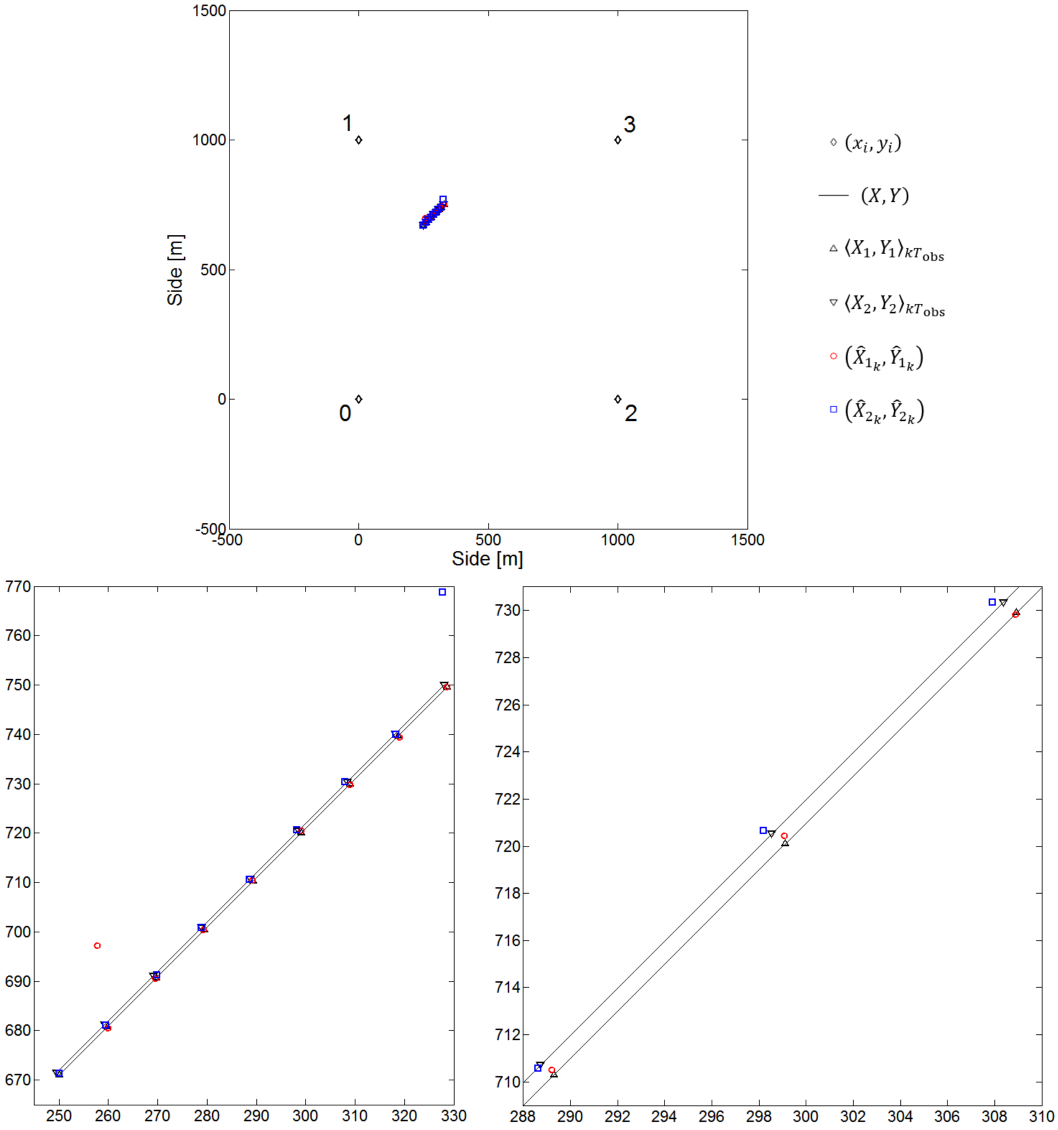

FIGURE 3.35: Scenario with two in-car jammers (the colored markers indicate the positions estimated in parallel by each EKF1).

the correct separation of the sawtooth waveforms: two jammers move at fast speeds along closely-spaced trajectories and their waveforms are modulated with comparable slopes. Tab. 3.4 complements the simulation parameters listed in Tab. 3.3. The scenario is illustrated in Fig. 3.35, where the jammers move at 100 km/h in opposite directions and pass each other at about 1 m of distance. The proximity of their locations with respect to the distances from the receivers is not a coincidence, because it is meant to stress the diversity in the power levels received by the four DoOs. More specifically, the evolution of the JNRs during this motion is plotted in Fig. 3.36: the received power levels allow us to distinguish between a "strong" jammer (i.e., $24\,\mathrm{dB} < \rho_1 < 33\,\mathrm{dB}$) and a "weak" one (i.e., $5\,\mathrm{dB} < \rho_2 < 14\,\mathrm{dB}$). The area between the DoOs is a square of side 1-km wide. The test results in Fig. 3.37 exhibit high accuracy for the strong jammer, which is tracked with errors under 1 m on the position and 10 km/h on the velocity. Likewise, the weak jammer is also located with sub-meter errors, whereas its velocity estimation is impaired by the stronger waveform and increase up to 20-30 km/h. The estimated locations are pinpointed in Fig. 3.35.

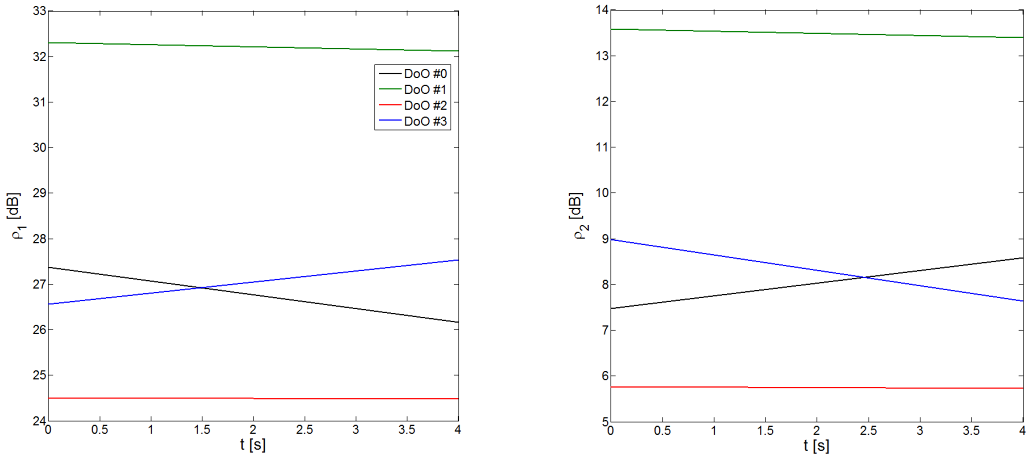

(A) Strong jammer with $P_t = 0\,\text{dBW}$.

(B) Weak jammer with $P_t = -20\,\text{dBW}$.

FIGURE 3.36: JNR levels received by the four DoOs.

3.4.3 Conclusions

recognizing the signatures of superimposed trains of linear chirps onto the TF domain Such sawtooth patterns typify the jamming waveforms generated by widespread commercial devices.

The cooperative and two-step estimation of position and velocity of a source of interference is extended to the case of multiple in-car jammers. By taking advantage of characteristic periodicities of jamming waveforms, a divide-and-conquer strategy is devised to tackle the sources of the attack: the waveforms are recognized in the TF representation and reconstructed through successive cancellations to produce separate CAFs, which are then used to separately estimate the TDOA/FDOA observables for every jammer. Since this operation depends on the success of the characterization algorithm introduced in the precedent section, the case study that evaluates the tracking performance is designed as a proof of concept under conditions that ensure the correct distinction of the jamming waveforms. When the jammers are successfully characterized, indeed, an array of KFs may be assigned to track their positions and velocities in parallel and is fed with the TDOA/FDOA measurements accordingly. The assignment works by associating one tracker to every peak of periodicity that is identified in the new domain constructed with the ST. Notably, the availability of more DoOs is exploited to cross-check the relevant periodicities extracted from the snapshots received at the same time. Since the overall periodicity map is more robust to missed detections and false alarms, this boosts both both the characterization and the tracking performance. For instance, the final localization accuracy in the scenario studied is below the meter over a 1-km square area even with a gap of 20 dB in the JNRs. Therefore, the presented method has the potential to track position and velocities of closely-spaced jammers with diverse transmit powers.

Whenever the overlapping sawtooths contained in the snapshots are many and similar, some failures could arise in the characterization algorithm and so affect the trackers. If one periodicity peak is mistaken for a jammer, a tracker will be erroneously assigned to a source that does not exist. Conversely, if one jammer is undetected because visible from one receiver only, then it goes missing. Over time, it is expected that these issues could be handled by scaling the array of KFs according to the consistency of their estimates. Anyhow, the robustness of the method proposed might benefit from a further layer of observation. More specifically, significant differences in terms of JNRs could be exploited to discriminate in terms of strength the interference received from different sources. Thereupon, a few comments are made in the previous section about the possible addition of a power adaptation at the stage of TF analysis with the ST. Despite the additional resilience granted by the power control, under the unlucky circumstance of two identical jammers placed closely and radiating the same amount of power, only one source of interference out of the two would be likely recognized and tracked. Nevertheless, in reality, even tracking just one jammer at a time would still enable to find and neutralize each threat sequentially, when a prompt action is taken. The same is hardly possible when applying state-of-the-art techniques for interference localization, since reliable time measurements cannot be normally taken under multiple jamming attacks.

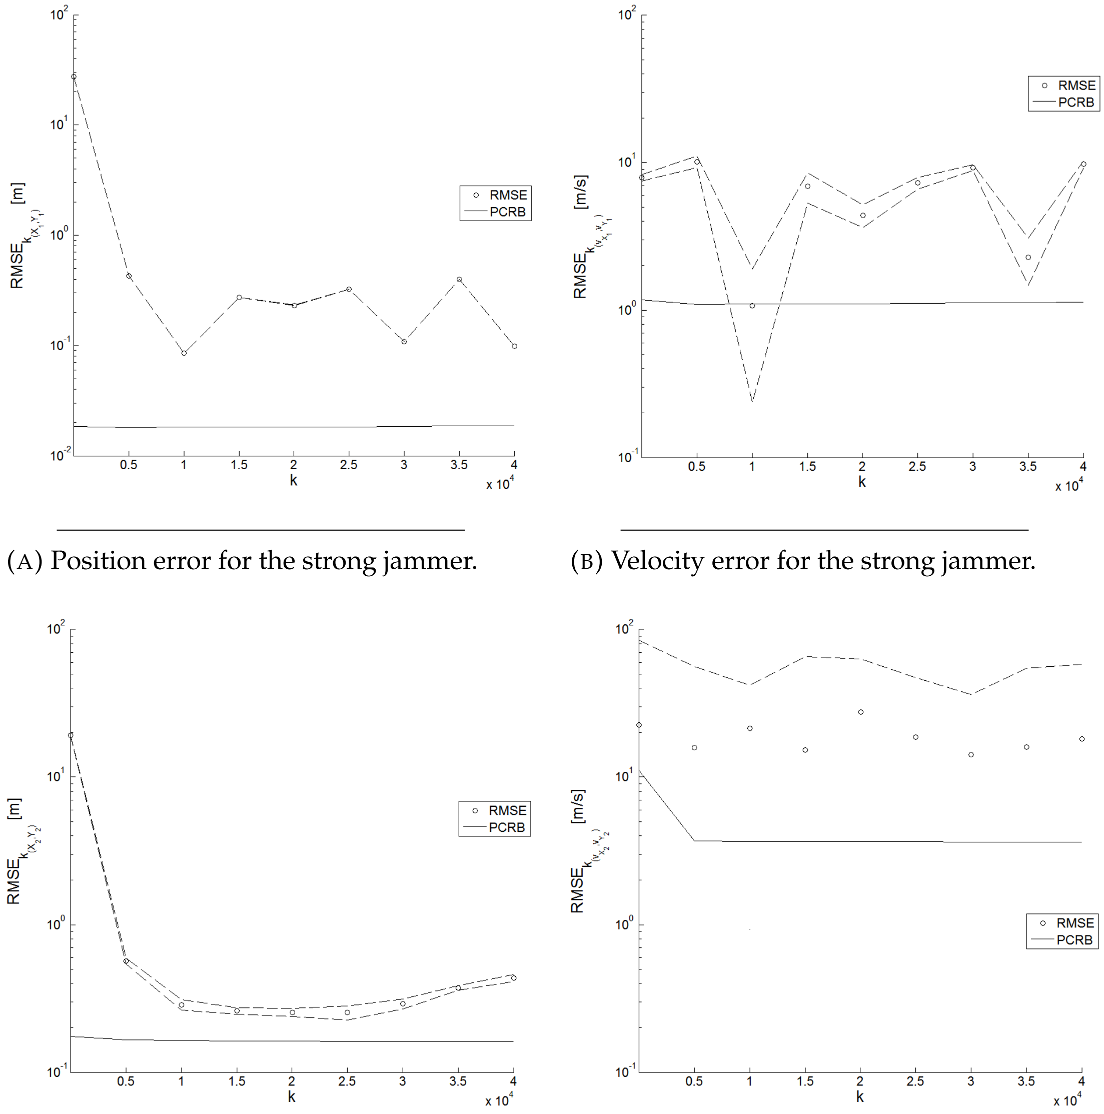

(A) Position error for the strong jammer.

(B) Velocity error for the strong jammer.

(C) Position error for the weak jammer.

(D) Velocity error for the weak jammer.

FIGURE 3.37: Performance in terms of RMSE of the estimation performed by each of the EKF1 in tracking both the position and velocity of the strong and weak jammers (the dashed lines trace the 95th percentile of the simulation precision).

Chapter 4

Received Interference Mitigation

The integrity and availability of GNSSs have gained great interest in the recent years, because of the ever-growing number of applications relying on an accurate PVT solution. All these systems essentially adopt a the DS-SS scheme: the navigation bits are modulated by a faster PRN code that spreads the spectrum in transmission over a bandwidth much larger than the minimum one imposed by the Nyquist criterion. This expedient benefits from an intrinsic immunity to in-band interference, because drops of the input SNR in reception can be compensated by the processing gain provided by despreading. Despite the degree of resilience of the DS-SS modulation, GNSS receivers are notably vulnerable to interference, because the signals transmitted by satellites are very weak on ground, namely 20-30 dB below the thermal noise floor. At the output of correlation with the PRN code, the interference spectrum has then the form of "extra noise", which deteriorates the carrier-to-noise-density power ratio (C/N_0), as shown in [165]. The consequences then range from the degradation of the accuracy and the integrity of GNSS navigation message to the disruption of the availability and continuity of the PVT solution, when the processing gain is not sufficient to recover the data. When the outage lasts for seconds, it results in a denial of service. Depending on the interference power spectral density within the bandwidth after despreading, a receiver may encounter conditions of increasing seriousness.

1. Low noise compromises the accuracy of the pseudoranges and causes sporadic losses of lock into the carrier and frequency tracking loops as well as the missed acquisitions of the weakly visible satellites. In other words, the channels with low C/N_0s are the first to go off or unstable.

2. Moderate noise worsens the previous effects on more channels. When few satellites can be acquired and track, the poorer geometry of the serving constellation has an important impact on the final positioning and timing accuracy. The integrity of the message decoded is compromised as well, because more errors are introduced, stressing the consistency checks.

3. High noise due to the strong interference provokes a near-far problem, so that the weak satellite signal is compressed by the AGC stage to a fraction of the dynamic range in order avoid the saturation of the ADC. Given the limited resolution in bits of the quantizer inside consumer-grade receivers (e.g., 1/2 bits), few tens of decibels in terms of JNR are enough to totally obscure the reception of the useful signal. When fewer than four satellites can be acquired and tracked, no PVT solution is possible and the denial of service is irreparable. Less precise and more robust fall-back systems (e.g., based on the cellular modem) should be then up and running.

The reader may find some studies dedicated to the impact of jamming attempts onto GNSS receivers in [165, 166] and others.
For safety-critical, mission-critical, and business-critical applications dependent on the GNSSs, the effects listed above easily prevent them from being trustworthy or stable as required by specifications. Furthermore, these consequences could resemble that of attenuation and multipath in dense urban areas. Therefore, the detection of the jamming attempts is the first necessary step to allow for countermeasures. Hereinafter, the presence of jammers is assumed to be known and the focus is on the receiver-side mitigation of interference.

4.1 State of the Art

Interference detection and mitigation are mechanisms aimed at ensuring the quality of service that is provided by the GNSS receiver and requested by positioning and timing applications. The detrimental effects of jammers provide also plenty of tangible clues to reveal their presence based on pre-correlation or post-correlation observables, like the AGC gain or the channel C/N_0, respectively. In this regard, relevant examples of methods for jamming detection are published in [167–176] and other papers. Besides the very powerful sources of interference built for criminal, terrorist, and military purposes, a variety of portable and inexpensive civilian jammers are available on the market that is accessible to the public. Their in-car versions are particularly popular. As mentioned in Chapter 3 and assessed in [146], the vast majority of these devices rapidly modulate in frequency of tone to obtain a wideband chirp-like waveform, which is repeated with a sawtooth pattern. By doing so, they can maximize the ratio between the average transmitted energy and the peak power. They typically feature power up to 1 W, average bandwidths of 15 MHz or more per channel, and repetition periods of about 10 µs. The Finnish Geodetic Institute analyzed the effects of these jammers on various consumer-grade GPS receivers working both in single-band and dual-band mode (i.e., u-blox, Nokia, and NovAtel) and published the results in [177]. A GPS L1 jammer of 13-dBm nominal power and 14-USD list price and a GPS L2/L5 33-dBm jammer worth 130 USD were employed. The experiments showed that the horizontal positioning errors dramatically increase both in their mean and variance and that the PVT solution availability drops from the 100% to just 8-26% in presence of interference, depending on the receiver as well as the on the jamming-to-signal power ratio. In addition, the experimental outcomes of [146] demonstrated that jammers as such can impair both acquisition and tracking performance of consumer-grade receivers in a range up to 9 km, approximately. This range far exceeds that advertised by the retailers, which rather erroneously indicate the maximum distance for effectively interfering strong wireless communications. Therefore, equipping GNSS receivers with anti-jamming modules is a crucial upgrade to guarantee the reliability of the PVT solution. A comprehensive review of the state of the art has been recently published in [178]. The demand for cost-effective implementations has motivated research on digital signal processing techniques for the excision of received jamming waveforms, which avoid the addition of expensive antenna arrays or inertial measurement units. In other words, under the assumption of constrained hardware, we consider here only software techniques: they process the raw I/Q samples at the output of the ADC/AGC loop in a domain where the powerful chirp-like interference exhibits distinguishable characteristics with respect to the weak GNSS signals, which in turn are dominated by noise. These mitigation methods may be classified based on the processing domain according to [179].

- Time-domain techniques need no transformed domain information. In this context, adaptive notch filtering that makes use of FIR or infinite impulse response (IIR) filters is proposed in [179–184]. The self-correlation among the filtered samples is negligible, as it does not hinders the relatively longer PRN code sequence. This approach is effective for interference mitigation as long as the jamming waveform is modulated with either constant or slowly time-varying IF, so that the set up and adaptation of the notches can keep up with them. In fact, they do not cope well with the fast sawtooth modulations of an unknown number of in-car jammers. In this context, an aid may come from the use of a KF to track quick frequency variations, as done in [131]. Besides being practically limited to mono-component interference, this addition, however, is sensitive to any possible mismatch between the received waveform and the model underlying the estimation. Within this category also falls another method named pulse blanking (PB), which was recently suggested in [185] for excising fast-modulated chirps into narrowband GNSS receivers; for instance, using a TV tuner with 1 MHz at baseband as front end. This low-complexity cancellation capitalizes on the fact that when only a fraction of the jammed spectrum is captured by the receiver bandwidth, it will actually resemble a periodic sequence of wideband pulses.

- Frequency-domain techniques search for interference by analyzing the FFT of the received samples as in [186]. They are effective only when a relatively small number of spectral components are contaminated by interference, which happens if the waveform has a sparse energy spectral density. Being only instantly narrowband, the frequency-swept tones radiated by in-car jammers, nonetheless, easily occupy the whole receiver bandwidth, thus saturating the spectrum under analysis.

- Representations in the TF domain map nonstationary signals that are dynamic both in time and frequency, as we discussed in Chapter 2. By doing so, they have the potential to overcome the shortcomings of the previous rejection approaches, while sparing more energy for the useful satellite signals. Indeed, they can reveal and extract the time-varying spectral content that characterizes the sawtooth waveforms commonly transmitted by in-car jammers. For anti-jamming countermeasures, this capability is exploited to estimate the IF of mono-component interference, such as the conventional train of linear chirps, and then suppress the interference according to the estimation. The cancellation is performed either by adapting the time-varying notch filter of [43] or by constructing the subspace orthogonal to the one of the jammer and then projecting the incoming samples onto it, such as in [93, 187, 188]. Alternatively, in [189], a similar outcome is obtained by subtracting from the jammed signal the jamming waveform re-generated based on the amplitude and phase that are estimated in the respective TF representation. A forth and last variant finally consists in applying a binary excision matrix, which is just multiplied to the matrix of the TF representation. This approach was originally presented in [190]. All these techniques are based on either the short-time Fourier transform (STFT) or the Wigner-Ville distribution (WVD).

The TF analysis includes many mathematical tools to deal with highly nonstationary waveforms, such as the sawtooths of interest. The STFT is a simple and linear representation with flexible complexity. However, it is subject to a constant trade-off between temporal and spectral resolution, which is determined by the choice of the analysis window. The window can be ideally optimized to represent linear chirps of a certain rate following [41] or adapted over time according to a concentration measure defined in [190]. Anyway, there is no window that can perform well with simultaneous waveforms of different characteristics. When it comes to the rejection of multiple jamming attacks, this limitation is certainly a problem. A second well-established tool in the literature is the WVD. As anticipated in Chapter 2, this TF representation is not tied to the same resolution trade-off of the STFT because it locates the components through a cross-correlation operation instead of windowing. However, due to the bilinearity, significant cross terms smear the representation and become numerous for frequency-modulated and/or multi-component signals. In order to reduce the amplitudes of the cross terms, the WVD can be smoothed by a time window. The smoothing window suppresses the cross-correlation between the signal components that are sufficiently separated in time, but it also reduce the spectral resolution, thus indirectly reintroducing a trade-off. Another approach to improve the basic WVD assumes that the parametric form of the signal is known a priori. For instance, a Wigner-Hough transform is exploited in [180] to retrieve the IF of a linearly-modulated chirp compatible with the interference generated by jammers. The distribution is then integrated over parametric curves that have the same form as the waveform IF, then searching for peaks in the new TF domain. This approach is generalized in [191] for jamming waveforms of arbitrary and known forms. The down-side of this and any model-based technique, however, is that they provide reliable performance, as long as the model matches the reality. This is obviously impossible whenever various and diverse waveforms are overlapped. Otherwise, the Hadamard product between WVD-based distributions was employed in [176] to reduce the cross terms to a large extent, while preserving a decent TF resolution. Generally, the critical aspect shared by all anti-jamming modules based on TF analysis is the complexity, because interference mitigation is supposed to run in nearly real time, so that the incoming signal is cleaned little after the reception. As opposed to the WVD, the STFT can be built as a bank of parallel filters. This implementation is convenient as the filter delays are short and usually within the tolerations of latency constraints. A similar implementation is possible for the multi-resolution extension of this transform, the ST, and is formulated in [6] and Chapter 2. In the same paper cited, it also devised a TF sampling scheme for scaling the amount computations at the expense of the accuracy of the representation and any modified signal reconstructed from it. Indeed, the complexity of the original ST in [23] imply prohibitive computational and storage requirements. Formally, it is on the order of $O(N^2 log_2 N)$ operations and $O(N^2)$ storage units, which are necessary to process FFT of Eq. 2.10 for all frequency bins, namely for all the voices. The complexity is addressed through the following discussion in various and novel ways.

In the present chapter, the goal is to take advantage of the desirable properties of the ST for interference rejection. A few implementations of the forward transform are explored to accommodate concerns about feasibility. The ultimate purpose is, in fact, to achieve an anti-jamming module that is a *practical add-on* for GNSS receivers. The convenient feature that motivates the adoption of the ST is the capability of providing consistent TF representations, with little or no assumptions about the number and the type of waveforms received. Indeed, the progressive resolution trade-off underlying this transform is suitable to potentially concentrate at the output the energy of the dynamic spectral content under analysis, regardless of the input evolution in time and frequency. The sole assumption regards the relative power of the incoming signals. Any useful signal should be below the noise floor and undetectable in terms of energy (e.g., DS-SS communications and GNSSs), whereas the waveforms under analysis should be easily **distinguishable from the background noise**. In other words, for the target application, the jammed transmissions once at the receiver antenna should be dominated by noise with negative and low SNRs, while the overall JNR should be positive and sufficiently high to make the impact of jammers stand out against the noise. The same considerations are mentioned for the methods of jamming characterization and localization in Chapter 3. The work behind the next sections is collected by two publications in [3] and [7]. The rejection methods presented are **non-parametric**, in the sense that they do not rely on any a-priori model or knowledge of the signal. As such, they are actually applicable to any kind of modulation, as long as it is powerful enough. The effectiveness of the interference mitigation are evaluated by emulating a GNSS receiver under jamming attack.

It is useful to make one remark about the technology constraints imposed by the addition of any anti-jamming unit: a GNSS receiver might benefit from the enhanced robustness to jammers only if it is also provided with a quantizer of sufficient resolution. As already mentioned, low-end mass-market receivers are equipped with 1/2-bit ADCs. Besides being prone to saturation, their short dynamic range clearly frustrate the effectiveness of any interference mitigation technique based on digital signal processing. Therefore, hereinafter, we implicitly consider devices that have the good enough specifications to make TF analysis worthy. For instance, a generic setup might include an SDR board like those shown in Fig. 3.1 with 8-bit quantization or higher and have an attenuator to take care of the absence of AGC. Example of testbeds may be found in [182] and [192]. Furthermore, devices used for critical applications can have between 6 and 14 bits (e.g., in [144, 145]).

4.1.1 GNSS Receiver Operation in Brief

The following summary recaps the basic operation of a GNSS receiver and introduces the reader to the terminology used in the rest of the chapter. More in-depth details about receiver technology may be found in [193].

The conventional architecture of a receiver is made of two steps that leverage on the structure of GNSS signals. After being down-converted to baseband or to an intermediate frequency at the front end, the received signal is still modulated by a ranging code and a much slower stream navigation bits. Ranging codes are quasi-orthogonal PRN sequences used to multiplex several satellite signals onto the same frequency band, according to the direct-sequence code division multiple access (CDMA). Being modulated with the code, the information spectrum is spread over a bandwidth much larger than the one actually necessary. This transmission technique is known as DS-SS and provides a degree of resistance to narrowband interference and interceptions. At the acquisition circuit, the procedure is reversed: the receiver searches for visible satellites by repeatedly cross-correlating the received signal with local replicas that embed all possible PRN sequences known a priori. Every replica is generated at the carrier frequency for a grid of code delays and Doppler frequency shifts. If the satellite associated to the searched sequence is visible, then the resultant AF exhibits one sharp correlation peak of magnitude, as shown in Fig. 4.1. The coordinates of this peak provide coarse estimates of the code phase and carrier frequency pair that are then used to initialize parallel tracking loops, which continuously refine them. Therein, the signals acquired are tracked in parallel by dedicated channels. As an alternative to the Doppler shift, the carrier phase is a more precise but ambiguous measurement, because of the higher temporal resolution provided

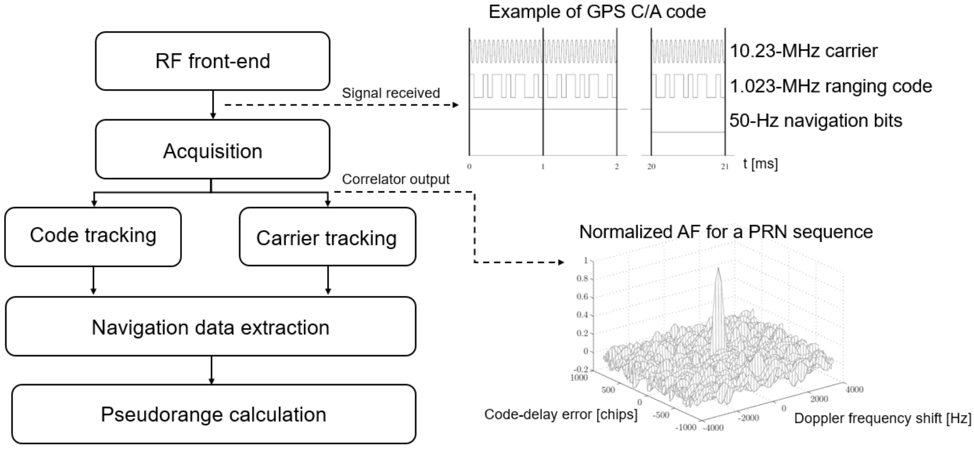

FIGURE 4.1: Conventional architecture of a GNSS receiver.

by the short wavelength. Within each enabled channel, as the local replica gets synchronized with the incoming signal, the carrier wave and the ranging code are demodulated. Next, the navigation message is decoded. The data so extracted and corrected contain the ephemerides to calculate the current position and velocity of satellites along their orbits, the bias parameters to account for the slow drift of the atomic clocks in space, an almanac collecting a history of recent ephemerides, and other complementary information. The navigation data are needed by the receiver to achieve final accuracies on the order of meters, to check the health of satellites, and to speed up the time to fix (i.e., warm/hot start). Together with the code and carrier measurements, they enable the precise estimation of the time of arrival and, so, of the pseudorange. If four or more satellites are acquired and tracked, the timely resolution of the trilateration equations returns the current PVT of the receiver.

4.2 Frequency-Adaptive S-Transform

Chronologically, the frequency-adaptive ST (FAST) was investigated before the complexity-scalable ST in Chapter 2. The idea here is to reduce the computational burden of the ST by computing the transform only for certain frequencies. According to our approach, these frequencies are the ones concentrating most of the energy of the signal under analysis. In order to identify them, a detection is performed onto the power spectrum by computing the FFT before the FAST. The reduction of complexity depends on the input power spectral density: if it exceeds the threshold over all the bandwidth, then the output of the FAST coincides with the fully-redundant ST. As opposed to the complexity scalable ST, therefore, the computational efficiency of the FAST is signal-dependant.

To introduce the FAST, it is useful to recall the discussion in Chapter 2. The ST is a TF analysis tool that improves on the STFT by adopting frequency-dependent Gaussian windows in order to provide progressive TF resolution. Windows of wide deviation represent the slowly-modulated components with high spectral resolution, while fast-modulated components are located with higher temporal resolution through windows of short deviation. The discrete-time ST of a sequence of N digital samples $x[n]$ is a complex-valued square matrix of order N defined by

$$\mathrm{ST}_x[n,p] = \sum_{m=0}^{N-1} x[m]w[n-m,p]\mathrm{e}^{-j\frac{2\pi}{N}pm} \tag{4.1}$$

for $n = 0,...,N-1$ and $p = -N/2,...N/2-1$, where the Gaussian windows $w[n,p]$ are defined in Eqs. 2.6 and 2.7 with standard deviations inversely proportional to the frequency index p. This matrix can be implemented as N DFTs (i.e., FFTs) according to Eq. 2.10. This formulation of the forward computation points out the direct relation between the ST and the FT: the amplitude and phase of local components in the two-dimensional TF domain are globally collected into the one-dimensional Fourier spectrum. In the absence of intermediate modifications, the DFT of the input $x[n]$ denoted by $X[p]$ can be exactly retrieved through the aforementioned frequency inverse (FI) as the output of the time-averaged representation:

$$X\left[p\right] = \sum_{m=0}^{N-1} \mathrm{ST}_x\left[m,p\right] \tag{4.2}$$

so that we have

$$x\left[n\right] = \mathrm{DFT}_q^{-1}\{X[q]\} = \frac{1}{N}\sum_{q=-N/2}^{N/2-1} X\left[q\right]\cdot e^{j2\pi\frac{q}{N}m} \tag{4.3}$$

with one inverse FFT (IFFT) block. The idea of the FAST proposed in [3] comes from the implications of Eq. 4.3. In brief, if one is interested in analyzing the TF representation only at the frequency bins that concentrate most of incoming energy, the forward computation of the ST could be adjusted accordingly. Thence, the ST is processed only for the voices where the power spectral density rises above the noise power level. In other words, the TF analysis is restricted to the spectral components that contain meaningful energy. The resultant computational efficiency depends on the signal power spectrum. The FAST is a suitable tool for mitigating the received interference in a few steps.

1. First, a *preliminary* detection stage identifies the frequency bins where the input power spectral density exceeds the noise level, because they contain TF components compromised by the interference. The bins apparently unaffected are forwarded to the last step.

2. The computation of the ST is restricted to the voices that correspond to frequency bins identified. The computational burden of the TF representation in then tailored to the visible interference power spectrum.

3. A TF detection stage determines the timings of the TF components affected by interference in the ST. They are detected at the time instants where the energy rises above the noise floor.

4. Once detected in time and frequency, the interference components are blanked by directly multiplying the ST with the negative of the decision binary mask. This process essentially filters out solely the significant energy in the TF domain. The principle behind masking the components most interfered though the ST is clear from Fig 4.2 and differs from conventional techniques, which remove the interference either by blanking the whole spectrum in certain time instants or by filtering certain frequency bands over the whole time.

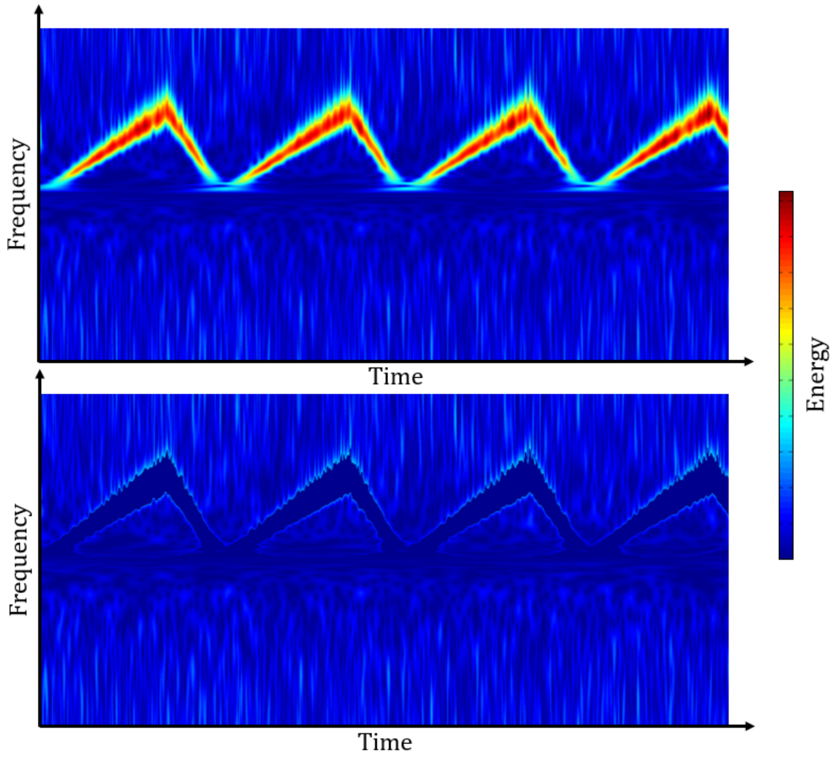

FIGURE 4.2: Energy of a jamming sawtooth waveform represented by the ST before and after the TF components containing significant interference power are detected and blanked (eventually an interference-free signal is supposed to be recoverable by inverting the transform with the FI).

5. After the blanking, the TF components in the masked ST are time-averaged to recover the corresponding bins of the clean Fourier spectrum without interference power. These bins are combined with those forwarded from the first step and processed all together by a final IFFT block. The time signal at the output of the FI is supposedly interference-free.

This approach is not model-driven, works with batches of samples and has a forward complexity equal to that of the FFT times the number of voices identified at the preliminary detection stage. The FI is then very efficient. Moreover, filtering the interference by simply blanking the respective TF representation replaces the need for an array of adaptive notch filters. The steps for rejecting the interference through the FAST are summarized in Fig. 4.3.

It is important to stress on the fact that the multi-resolution ST is more advantageous than the fixed-resolution STFT when the signal analysed is unknown; in the application or interest, no assumption is necessary about the number and the TF features of the interference to mitigate. With the procedure described above, any kind of frequency-modulated jamming waveform can be automatically removed by blanking the respective energy in the TF representation: continuous waves, sawtooths, pulses, frequency-hopping, etc. Obviously, if the interference power spectral density covers extensively the TF plane in reception, such as in the case of wideband noise-like interference, masking the unwanted energy at the input would not provide any benefit, because the useful signal is entirely wiped out as well. In reality, a jamming attack as such is hardly feasible from distance, since it requires a huge

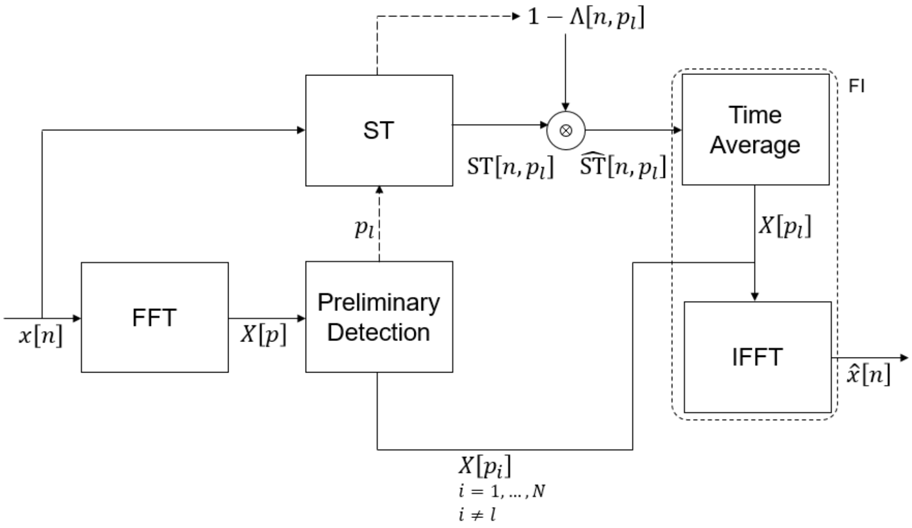

FIGURE 4.3: Block diagram of the interference rejection stages based on the FAST.

amount of power. Anyway, this unlucky case is considered as irredeemable with digital signal processing techniques, like the saturation of the front end.

4.2.1 Preliminary Detection Stage

The GNSS signals have noise-like characteristics, because they are typically buried under the flat noise spectral density. On the contrary, the power spectral densities of common jamming sawtooth waveforms have usually peaks, the intensity of which depends on the product between the bandwidth and the repetition period of the jammer. An example is shown in Fig. 4.4. Since the denial-of-service attempt tries to overpower the useful signals in reception, it is also intentionally powerful with respect to the noise. To exploit this fact, the first step is to identify the portion of the incoming spectrum that concentrate most of the jamming power. This task is solved as a binary decision problem by hypothesis testing: if the power level of a certain frequency bin exceeds a certain threshold, it is associated to interference and flagged for further processing through the next steps. Let us consider additive white Gaussian noise from this moment on. The threshold denoted by λ^{FT} can be simply pre-determined with a fixed false-alarm rate $P_{\mathrm{fa}}^{\mathrm{FT}}$ under the hypothesis H_0 of interference absence according to

$$\lambda^{\mathrm{FT}} = -2\sigma_n^2 N \ln\left(P_{\mathrm{fa}}^{\mathrm{FT}}\right) \tag{4.4}$$

with

$$P_{\mathrm{fa}}^{\mathrm{FT}} = \mathrm{prob}\left(|\mathrm{X}\,[p]|^2 > \lambda^{\mathrm{FT}} \middle| \mathrm{H}_0\right). \tag{4.5}$$

Whenever the received interference power is comparable to the noise level, the decision test is prone to missed detections. However, under this circumstance, the receiver-side processing gain after despreading is expected to be sufficient to decode the navigation message without difficulties.

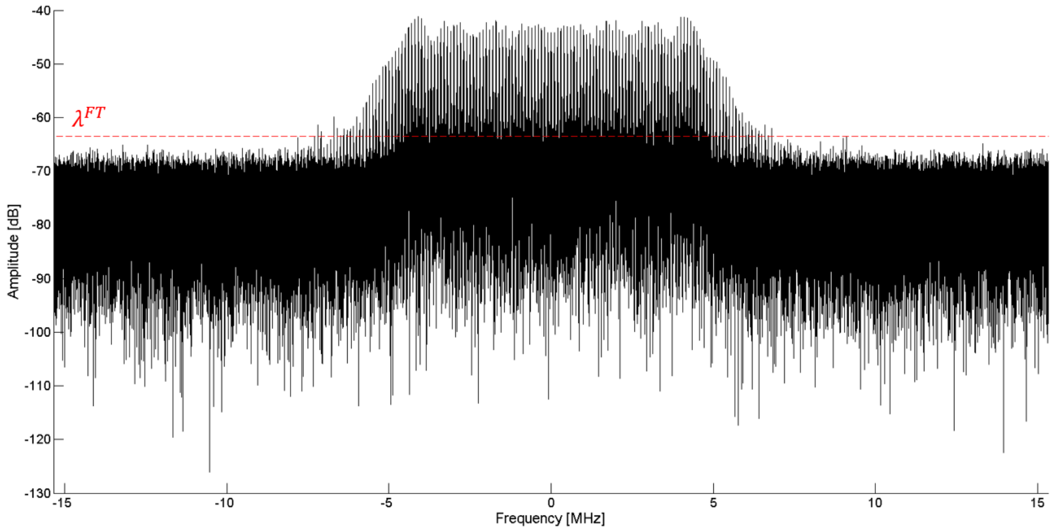

FIGURE 4.4: Baseband GNSS signal overpowered by a 10-MHz sawtooth waveform (the complexity of the FAST is reduced by a factor $N/L = 13.82$ compared to the original ST).

4.2.2 Time-Frequency Detection Stage

The preliminary detection stage identifies the jammed portions of the Fourier spectrum. However, the time instants at which the frequency bins are actually exposed to the received interference are still unknown. Given the instantly narrowband nature of the sawtooth patterns used to modulate jamming waveforms, at each frequency only a few samples are periodically contaminated by interference. Consequently, removing the whole power identified at the preliminary stage would result in an unnecessary loss of the energy carried by the GNSS signal hidden in noise. Such a limitation is overcome by resorting to a second detection stage, which is performed in the TF domain. The ST is then computed only for a number L of voices at the indices p_l with $l = 1, ..., L$, which are flagged at the first stage, thus sparing computational power and memory resources. As a result, the complexity is reduced by the factor $N/L \geq 1$ compared to the complete ST. Over the voices selectively computed, the second decision test compares the energy of every TF component with a frequency-dependent threshold $\lambda^{\text{ST}}[p_l]$ pre-determined at constant false-alarm rate. Under the hypothesis H_0 of interference absence, the probability of false alarm $P_{\text{fa}}^{ST}[p_l]$ for a specific voice of the ST is obtained by generalizing the formulation given in [43] for the spectrogram. It is so defined as

$$P_{\text{fa}}^{\text{ST}}[p_l] = \text{prob}\left(|\text{ST}[n,p_l]|^2 > \lambda^{\text{ST}}[p_l]\right) = \exp\left(-\frac{\lambda^{\text{ST}}[p_l]}{2\sigma_n^2 E_w[p_l]}\right) \tag{4.6}$$

that is equivalent to Eq. 3.73, where $E_w[p_l]$ is the energy of the Gaussian window $w[n, p_l]$. As anticipated, the false-alarm rate may be fixed to $\overline{P_{\text{fa}}^{\text{ST}}}$ following the rule of thumb in Eq. 3.74. The resultant threshold is

$$\lambda^{\text{ST}}[p_l] = -2\sigma_n^2 E_w[p_l] \ln \overline{P_{\text{fa}}^{\text{ST}}}. \tag{4.7}$$

4.2.3 Interference Masking

After identifying the jammed components in both time and frequency, the binary mask $\mathbf{\Lambda}$ returned by the binary decision test is used to blank the components detected. The elements in the matrix $^{\mathrm{ST}}$ are ones where the energy exceeds $\lambda^{\mathrm{ST}}[p_l]$. Therefore, the interference is simply filtered out through the following Hadamard product

$$\widehat{\mathrm{ST}}[n, p_l] = \mathrm{ST}[n, p_l] \otimes (1 - \Lambda[n, p_l]). \tag{4.8}$$

that is just an element-wise product among matrices. The masked representation in $\widehat{\mathrm{ST}}[n, p_l]$ is time-averaged into the spectrum $\hat{\mathrm{X}}[p_l]$, which is combined with the values at the frequency bins that went undetected through the preliminary stage. The final time signal $\hat{x}[n]$ is supposed to carry a negligible amount of residual interference power.

4.2.4 Case Study: Galileo Signal Acquisition

The method proposed is tested to enhance the performance in terms of code acquisition of a Galileo receiver that undergoes a jamming attempt. The acquisition circuit adopts the non-coherent channel combining scheme in [194] together with a maximum searching strategy. This combination is arguably the simplest implementation, since it does not take advantage of the signal structure. The samples captured by the front end are first demodulated into I/Q branches. Second, the I and Q streams are separately correlated with the local replicas of the data and pilot channels. The two outputs of the two correlators are non-coherently integrated into the AF. An a-posteriori decision is taken on the maximum peak of magnitude in this function: the received Galileo signal is considered as successfully acquired if the peak exceeds a certain threshold with a code delay error less than half of the chip time. The threshold η is obtained by fixing and inverting the following false-alarm probability

$$P_{\mathrm{fa}}^{\mathrm{ACQ}} = \exp\left(-\frac{\eta N}{\sigma_n^2}\right)\left(1 + \frac{\eta N}{\sigma_n^2}\right). \tag{4.9}$$

Without loss of generality, we may neglect the quantization losses and consider a single channel, where only the right satellite is visible thanks to the code orthogonality. The noisy samples received incorporate both the useful Galileo E1 OS signal and one jamming sawtooth waveform, which is modeled as described in detail in Chapter 3. The performance of the acquisition circuit adopted are evaluated in terms of the rate of code detection denoted by P_{acq} for different $\mathrm{C/N_0}$ levels. For this purpose, Monte Carlo simulations are run with the parameters listed in Tab. 4.1 over 300 iterations. The performance without anti-jamming mechanisms is shown in Fig. 4.5. Fig. 4.6 highlights the improvements achieved by excising the interference through the FAST. From the results clearly emerges the higher robustness of the acquisition, even when the $\mathrm{C/N_0}$ is low. The high dependency of the acquisition rate on the jammer TF characteristics is reasonable, because with the repetition periods changes the power spectrum and so the resolution of the FAST. For the sake of comparison, we evaluate the performance for the same method in Fig. 4.3, but masking the jamming energy analysed through a fixed-resolution STFT instead of the multi-resolution ST. The TF representation is so computed with a sliding Hamming window, which has one chosen size. The resultant rate of code acquisition versus the window size denoted by M_w is depicted in Fig. 4.7. In this test, the $\mathrm{C/N_0}$ tested is set to 45 dBHz, at which in Fig. 4.5 the ST performs well for the three jammers tested. The results of the

TABLE 4.1: Parameters for the case study of the acquisition of a jammed Galileo E1 OS signal.

Description	Symbol	Value
Carrier frequency	f_c	1575.42 MHz
Code duration		4 ms
Sampling frequency	f_s	$30 \cdot 1.023$ MHz
Intermediate frequency		$5 \cdot 1.023$ MHz
C/N$_0$ level		$\{30, 34, 38, 42, 45\}$ dBHz
Front-end equivalent noise passband bandwidth	B_{eq}	10 MHz
Jammer passband bandwidth	B_{RF}	10 MHz
Jammer repetition period	T_0	$\{1, 10, 100\}$ µs
In-band JNR level	ρ	10 dB
Preliminary detection false-alarm probability	P_{fa}^{FT}	10^{-6}
TF detection false-alarm probability	P_{fa}^{ST}	10^{-6}
Acquisition false-alarm probability	P_{fa}^{ACQ}	10^{-3}

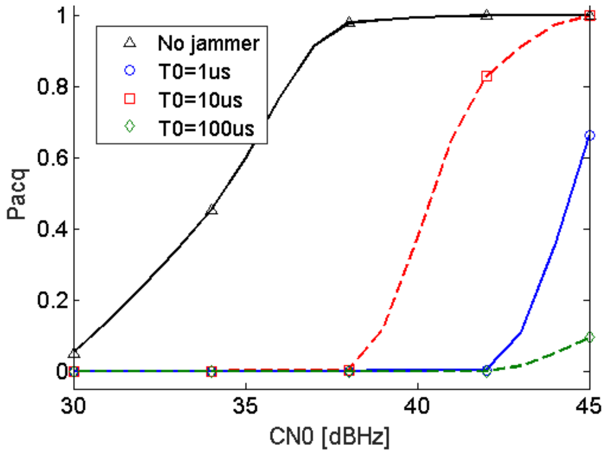

FIGURE 4.5: Code acquisition rate with and without interference (the black line is the nominal performance).

STFT instead confirm the expected dependency on the analysis window: none of the sizes is suitable to all the repetition periods. And if the representation is poor, any modification might do more harm than good, distorting and erroneously blanking useful energy. Ideally, the window shall rather be adapted to the specific TF characteristics of the waveform. However, in reality, the incoming interference could be made of many and unknown waveforms. Under this circumstance, the progressive resolution trade-off of the ST alleviate the critical role of the windows for the rejection effectiveness. As a result, multiple and diverse jamming attacks can be tackled, as demonstrated at the end of the chapter.

4.3 Time-Selective S-Transform

Along the FAST, we may likewise restrict the computation of the ST only for the time instants in which significant in-band energy is captured. This principle produces another variant of the transform, named time-selective ST (TSST), which is similar

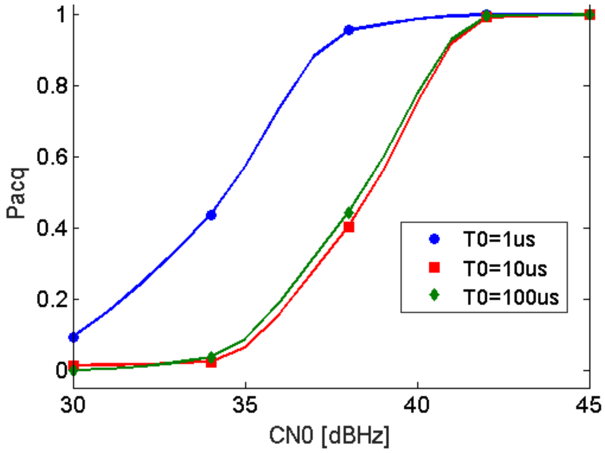

FIGURE 4.6: Code acquisition rate versus C/N_0 with interference rejection based on FAST.

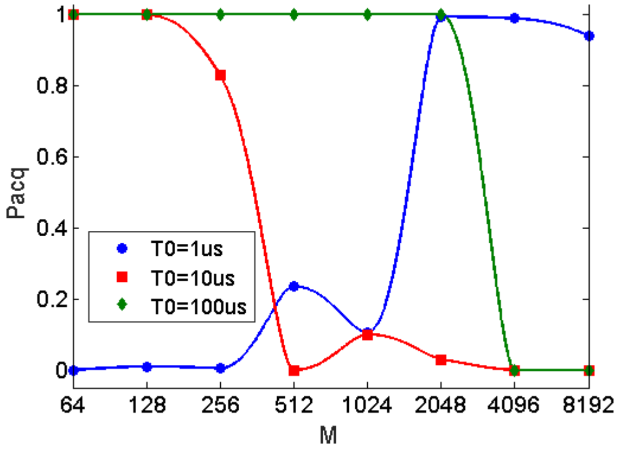

FIGURE 4.7: Code-acquisition rate versus the window size with interference rejection based on STFT and 45-dBHz C/N_0.

to the FAST. This one is particularly useful for the excision of jamming waveform in narrowband receivers. In fact, since jammers usually sweep a tone over a large bandwidth, the amount of interference power actually visible is likely affecting only short and periodic intervals of time in reception. The same principle has motivated the use of pulse blankers for anti-jamming modules: when the jammer bandwidth is much wider than the one of the front-end filter of to the GNSS receiver, then the incoming sawtooth is seen as a periodic sequence of pulses in reception. This fact is shown by the spectrogram in [185, Fig. 1]. Consequently, whenever the instantaneous power is suspiciously higher than noise, the corresponding sample can be simply blanked to remove the dominant interference components. This simple technique provides an immediate and effective layer of protection against interference with an extremely low-complexity implementation. The simplicity is the reason why it is a popular add-on for commercial receivers. The impact of blanking is equivalent to a shorter integration time at the acquisition circuit, while long cancellations can impair the stability of the tracking loops as well as introduce uncorrectable errors in the navigation message. For this reason, we overcome the shortcomings of pulse blanking in the framework of TF analysis. Alternatively to the approach offered by the FAST, the matrix of the ST can be computed only at the instants of index n_l where the magnitude of the digital signal $x[n]$ is above a certain threshold:

$$|x[n]| > \lambda^{\mathrm{PB}}. \tag{4.10}$$

As usual, the threshold is set as

$$\lambda^{\mathrm{PB}} = \sqrt{-2\sigma_n^2 \ln\left(P_{\mathrm{fa}}^{\mathrm{PB}}\right)} \tag{4.11}$$

at a constant false-alarm rate equal to

$$P_{\mathrm{fa}}^{\mathrm{PB}} = \mathrm{prob}\left(|x[n]| > \lambda^{\mathrm{PB}} \,\middle|\, \mathrm{H}_0\right) = \exp\left(-\frac{(\lambda^{\mathrm{PB}})^2}{2\sigma_n^2}\right) \tag{4.12}$$

in which the signal is hypothesized to be white Gaussian noise in the absence of jamming attempts. When it comes to highly quantized signals, the threshold can be set coherently with the histogram of the samples that are assumed free from interference. While the others are forwarded, the samples that satisfy Eq. 4.10 are searched for interference with the so-called TSST. Once the transform is computed for these time instants, the components apparently interfered are also identified in frequency as the ones that exceed the energy threshold of Eq. 4.7. They are finally filtered out from the TF representation by using the binary detection mask similarly to what done in Eq. 4.8. As opposed to the FAST, after interference blanking, the samples filtered through the TSST are recovered through the time inverse (TI), instead of the FI. They are finally recomposed with those initially forwarded and unfiltered to retrieve a time signal, which is supposed to be free from interference.

In conclusion, while the FAST performs a joint spectral and TF detection of the received interference, the TSST applies the same concept for a joint temporal and TF detection that essentially generalizes the PB. Both these transforms reduce the original amount of computations to obtain a TF representation with progressive resolution, but they differ by the consequent impact on the successive mitigation. Contrary to the complexity-scalable ST of Chapter 2, the efficiency is not arbitrarily chosen, since it is adjusted to the incoming energy spectral or temporal density.

4.4 Performance Assessment of Anti-Jamming Units

Given its advantageous properties of linearity and multi-resolution, the ST clearly emerges as a potential enabler for anti-jamming units, since it alleviates the need for a-priori knowledge about neither the number nor the characteristics of the jammers. This capability is especially useful when dealing with multiple overlapping waveforms, which cannot be analysed through the well-established STFT or WVD. In the following, we go through the preparatory results to be published in [7]. The focus is on the performance of a GNSS receiver equipped with a module for interference detection and mitigation based on differently-chosen portions of the same TF representation: the original and fully-redundant ST of [23], the FAST proposed in [3], and the novel TSST introduced in the precedent section. Their effectiveness is compared to that of the PB evaluated for one jammer in [185], which is here chosen as a reference. The PB is an appealing technique from the perspective of the implementation, because of its extremely low complexity.

TABLE 4.2: Parameters for the case study of multiple in-car jammers.

Description	Symbol	Value
Carrier frequency	f_c	1575.42 MHz
Code duration		1 ms
Sample record duration		36 s
Sampling frequency	f_s	$2.5 \cdot 2.048$ MHz
Front-end equivalent noise passband bandwidth	B_{eq}	2.048 MHz
Estimated noise power	$\hat{\sigma}_n$	20.54 dBW
Jammer passband bandwidths	B_{RF}	$\{5, 20\}$ MHz
Jammer repetition periods	T_0	$\{30, 10\}$ µs
In-band JNR level	ρ	$\{ 10, 20 \}$ dB
Sample batch	B	2048
Spectral pre-detection false-alarm probability	$P_{\text{fa}}^{\text{FT}}$	0.11
Temporal pre-detection false-alarm probability	$P_{\text{fa}}^{\text{PB}}$	0.11
TF detection false-alarm probability	$P_{\text{fa}}^{\text{ST}}$	B^{-2}

4.4.1 Receiver Setup

Our case study considers the reception of GPS L1 C/A signals. The 8-bit I/Q samples are re-generated with a GPS signal simulator (github.com/osqzss/gps-sdr-sim). The data used for the simulation have been collected in the receiver-independent exchange format (RINEX) by a LEICA SR9500, which is located at the GNSS reference station of the University of Bologna in Fig. 4.8. We post-process the signal through the steps listed below.

1. We emulate realistic jamming attacks coherently with the literature, such as [146], generating linear chirps modulated with the sawtooth patterns characterized in Tab. 4.2. A JNR level of 30 dB is already enough to nearly saturate the 8-bit receiver ADC in the absence of AGC.

2. We test different anti-jamming units that blank the incoming interference either in time or TF domains. They make use of PB, ST, FAST, or TSST. All of them rely on different detection statistics, which are nonetheless set with the equal false-alarm rate.

3. We assess the impact of the interference rejection through the whole processing chain of a GPS software receiver: acquisition, tracking, and PVT solution.

What is crucial but often neglected in the scientific literature is the evaluation of the effectiveness of the anti-jamming units across all the whole processing chain of a GNSS receiver. To pursue this task, we have customized the single-frequency GPS software receiver explained in detail in [193]. The receiver post-processes a record of I/Q samples, performing an initial acquisition and then tracking the visible satellites in a serial fashion. No adaptation of the quantizer and the gain is implemented. At least a full frame of 30 s is necessary in warm-start mode to obtain the ephemeris data, while a few more seconds shall be added to allow for the initial sub-frame synchronization. Despite its simplicity, other and similar software receivers are often used as a testbed in many publications in the field of GNSSs (e.g., in [185]). As far as our investigation is concerned, our customization of a GPS software receiver

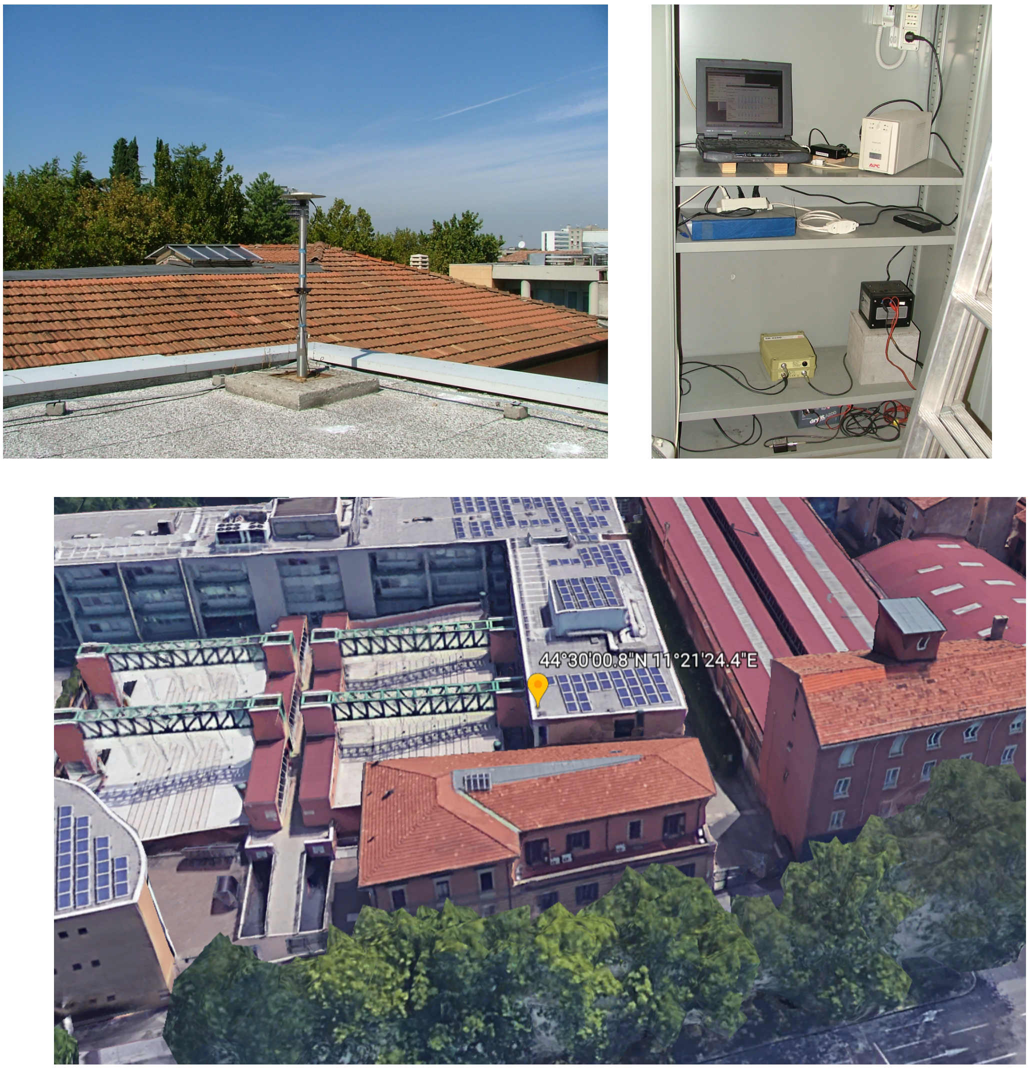

FIGURE 4.8: Pictures of the GNSS reference station of BOLG00ITA (epncb.oma.be) and respective geographical location in Google Earth.

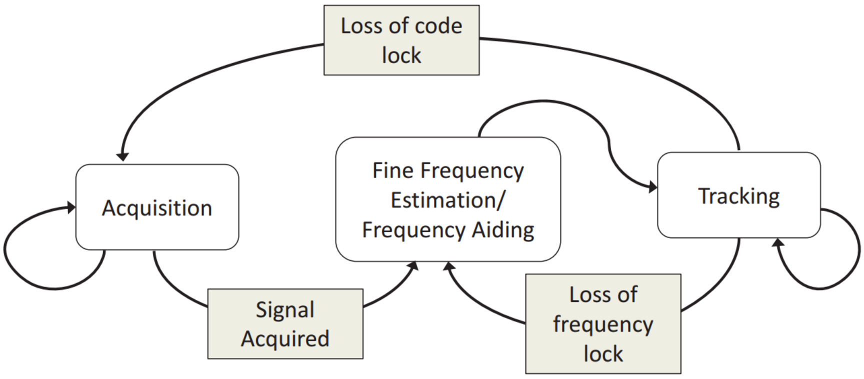

FIGURE 4.9: State machine implemented for every channel of the software receiver in [185].

provides a useful benchmark for the application under study. In fact, we are interested in assessing the performance variations, to which the reception is subject, depending on the methods implemented for interference mitigation. Achieving the best-possible robustness under jamming attacks is beyond of the scope of our study, as we focus on the enablers of effective interference rejection and not the architecture of the receiver. The incoming interference power is blanked by batch pre-processing the I/Q samples. This procedure emulates the process of an anti-jamming unit installed after the digitization and before the acquisition and tracking blocks, which operate separately. In this regard, it is worth mentioning that one major upgrade to the software receiver architecture was to enable the parallel execution of channels, which now can search for distinct PRN sequences independently. Furthermore, each channel was implemented as the finite state machine in Fig. 4.9 and [185, Fig. 4] with the aim of improving the resilience of their operation to the disruptions of the GNSS signals. The transitions to re-acquisition and frequency-aiding respond to possible losses of the code or carrier locks, respectively. The lock indicators monitoring the trackers are calculated as described in [195], based on estimates of the C/N_0 and the carrier-phase oscillation. The following performance metrics are extracted from the acquisition circuit, the tracking loops for the satellite with the highest acquisition metric, and the PVT calculation.

- The acquisition metric is computed as the ratio between the magnitudes of the maximum peak and the second-highest peak in the AF. If this metric is above the chosen threshold (i.e., 2.5), the satellite identified by the corresponding ranging code is assumed as acquired. The time elapsed before the first acquisition is also measured.

- The C/N_0 is estimated from the power levels of the I and Q samples at the outputs of the respective prompt correlators, after the code demodulation. In the following, this metric is shown only for the satellite that normally has the highest acquisition metric. During the time window in which the data were collected from BOLG00ITA, the satellite was clearly the one with PRN number (#) equal to 12, as proven in Fig. 4.10.

- Since the centimetre-level position of the receiver is known, the horizontal and overall (three-dimensional) error on the position solution can be calculated when there the navigation data are successfully recovered from at least four satellites. In normal conditions, the performance of the software receiver with the dataset tested achieves the accuracy in Fig. 4.11. Whenever the preamble is not found in the demodulated bits, the corresponding satellite is not used for the calculation of the PVT.

4.4.2 Preliminary Results

The numerical results of the tests performed with the setup receiver have two goals. On the one side, they demonstrate the susceptibility of PB to jamming waveforms with “slow” modulations as well as, by extension, to multiple and overlapping waveforms. The reason behind this weakness is intuitive: this technique tends to erase the most of the signal if the received interference power exceeds the amplitude threshold over large time spans. On the other side, they want to prove that the ST and its variants, namely the FAST and the TSST, can overcome the drawback of the PB and provide more stable performance to the receiver. Particularly, when the JNR is low and interference energy is persistently captured within the bandwidth of the front-end filter, the impact of the PB might do more harm than good: it could prevent the receiver from obtaining a fix that otherwise would be possible even without mitigating. Under the same circumstances, instead, the interference rejection based

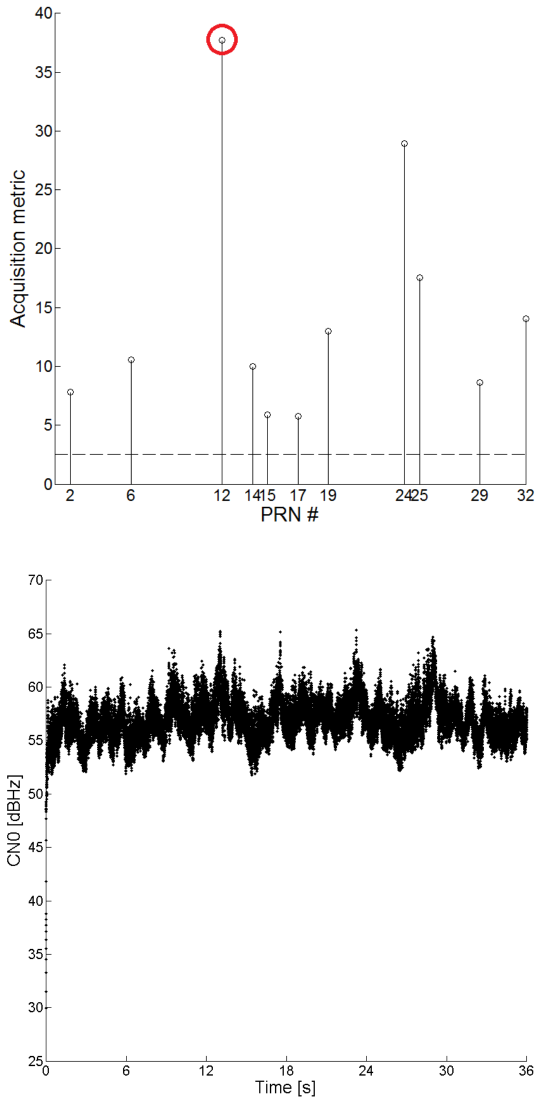

FIGURE 4.10: Normal acquisition and tracking performance for the reception of the GPS L1 C/A signal in the absence of interference.

on ST is not expected to degrade the performance, but possibly to improve them. The first test is carried out with a single jammer transmitting a slowly-modulated sawtooth waveform, which spans a 5-MHz bandwidth with periodicity of 30 µs. When the JNR is equal to 10 dB, the results in Fig. 4.12 demonstrate how the PB might be detrimental compared to the absence of interference mitigation. Indeed, despite the narrow band (i.e., 2 MHz) visible at the receiver end, such a waveform is not seen a sequence of pulses, because its energy lies within the front-end bandwidth for most of the time. Consequently, long portions of the signals are de facto removed by the PB with catastrophic consequences throughout the processing chain of the receiver. The weakness of this method presumably tends to be more severe if the avaialble bandwidth is larger than the one under test. On the contrary, the performance metrics obtained by using any of the STs exhibit some cautious enhancements in acquisition and tracking, without a meaningful effect on the final geolocation accuracy. The same test is also run with 20-dB JNR. The results in 4.13 verifies again the same detrimental repercussions of the PB, as opposed to the slightly beneficial

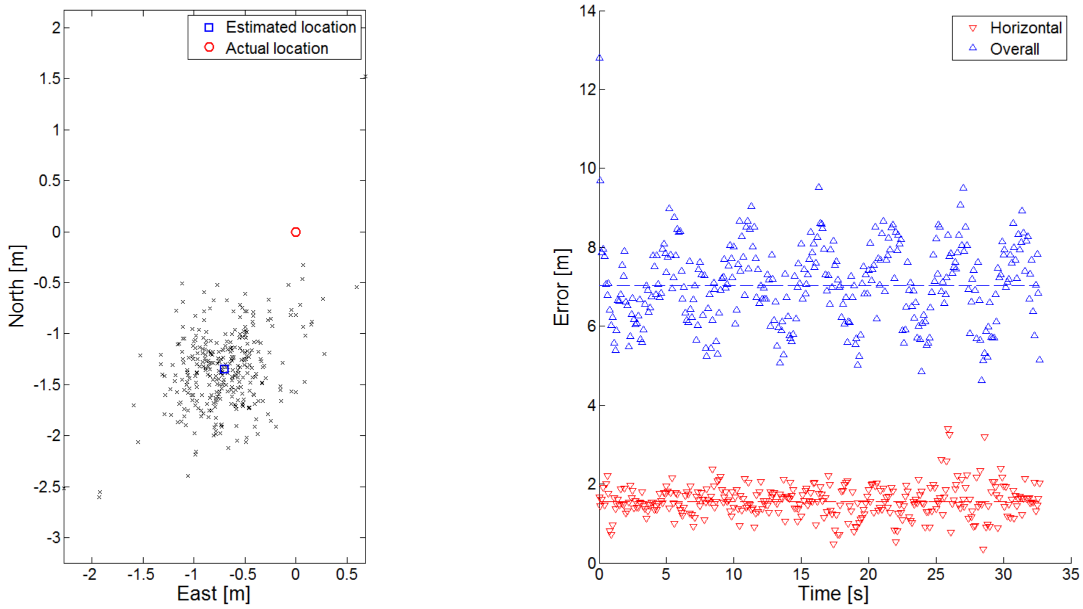

FIGURE 4.11: Geolocation accuracy of the GPS software receiver in the absence of interference with respect to the universal transverse mercator (UTM) coordinate system (the black markers are the positions estimated with partial data and indicate the precision around the final estimate).

impact especially of the ST. The apparent improvements in the acquisition and tracking performance are noticeable. Nevertheless, with our setup receiver, none of the methods lead to the calculation a useful navigation solution. In other words, too few channels are tracked for a fix, and the navigation data recovered is anyhow too corrupted. A second test overlaps the previous jamming waveform with a second sawtooth, which has a faster modulation of 20 MHz in 10 µs. The results for an overall JNR of 10 dB for the two jammers combined are shown in Fig. 4.14, where hold the same considerations made for Fig. 4.12.

4.5 Conclusions

Given the non-stationary nature of the waveforms typically transmitted by in-car jammers, there is an emerging research trend towards the application of TF analysis to anti-jamming modules. A higher degree of resilience is crucial in critical applications relying on the GNSS service on ground. In this context, the potential employment of the ST for interference detection and mitigation has been investigated. The latency and computational power necessary to process such a multi-resolution TF representation can be properly reduced by exploiting the specific temporal or spectral characteristics of the received interference. The adaptation is automatic and non-parametric, and it is provided with the novel versions of this transform, namely the FAST and the TSST. Besides, the complexity of the original ST might still be practically affordable when processing the incoming signal in short batches. The efficiency cost is anyhow motivated by the remarkable enhancement in terms of robustness that is granted by the adoption of the ST. Indeed, in comparison to the commonly-used PB, the jamming rejection methods built around this transform and its variations minimize the receiver susceptibility to interference, without any knowledge of the number and characteristics of the sources. Generally, thanks to progressive resolution trade-off, the unwanted energy is well identified in time and frequency and blanked therein. The tests made with a GPS software receiver probe the potential behind the use of the ST, the FAST, or the TSST as moderate-complexity add-ons to GNSS receivers. Depending on the resolution that can be achieved by the receiver specifications, they could reject interference attempts from many and diverse jammers. The same methods are also directly applicable to protect DS-SS communications, as long as the useful signal has a flat power spectral density or is buried in noise.

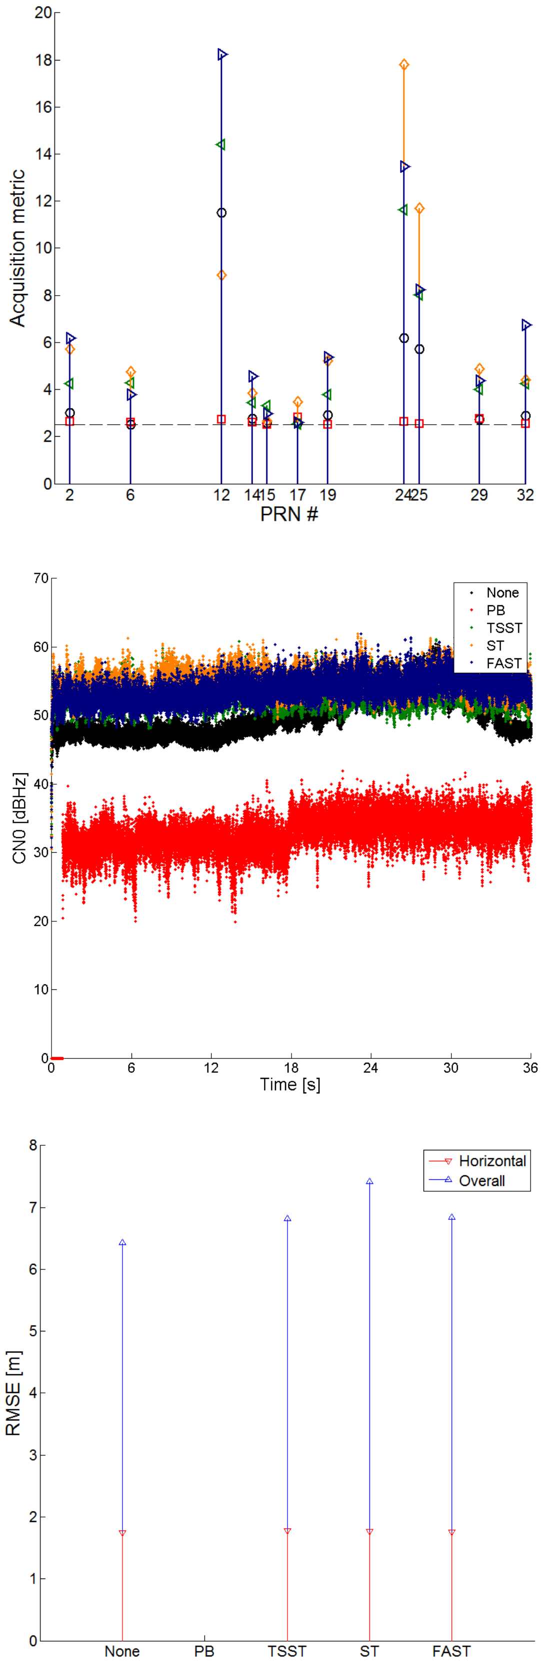

FIGURE 4.12: Acquisition, tracking, and positioning performance in the presence of one jammer with 10dB of JNR.

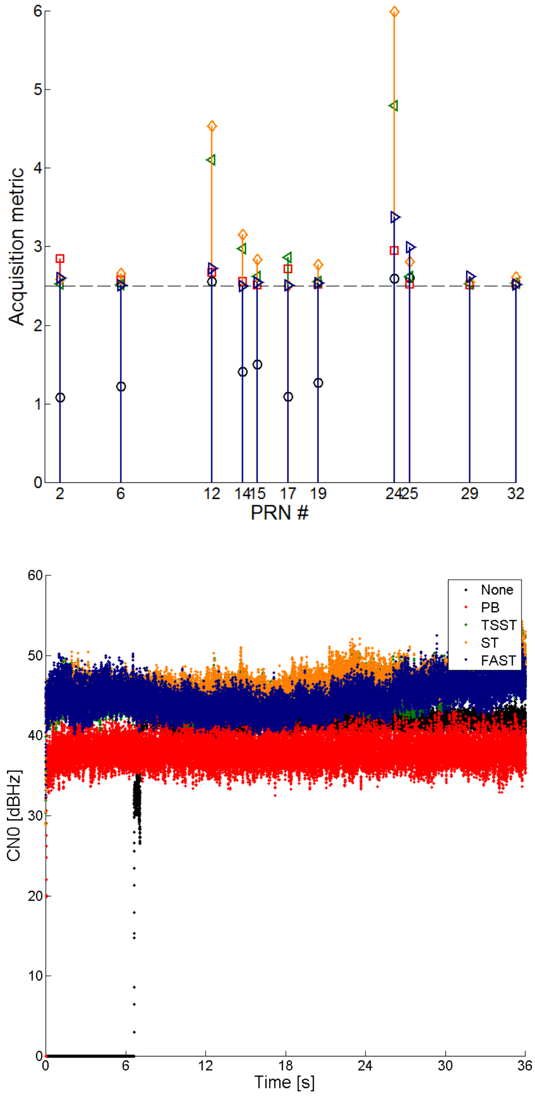

FIGURE 4.13: Acquisition and tracking performance in the presence of one jammer with 20dB of JNR.

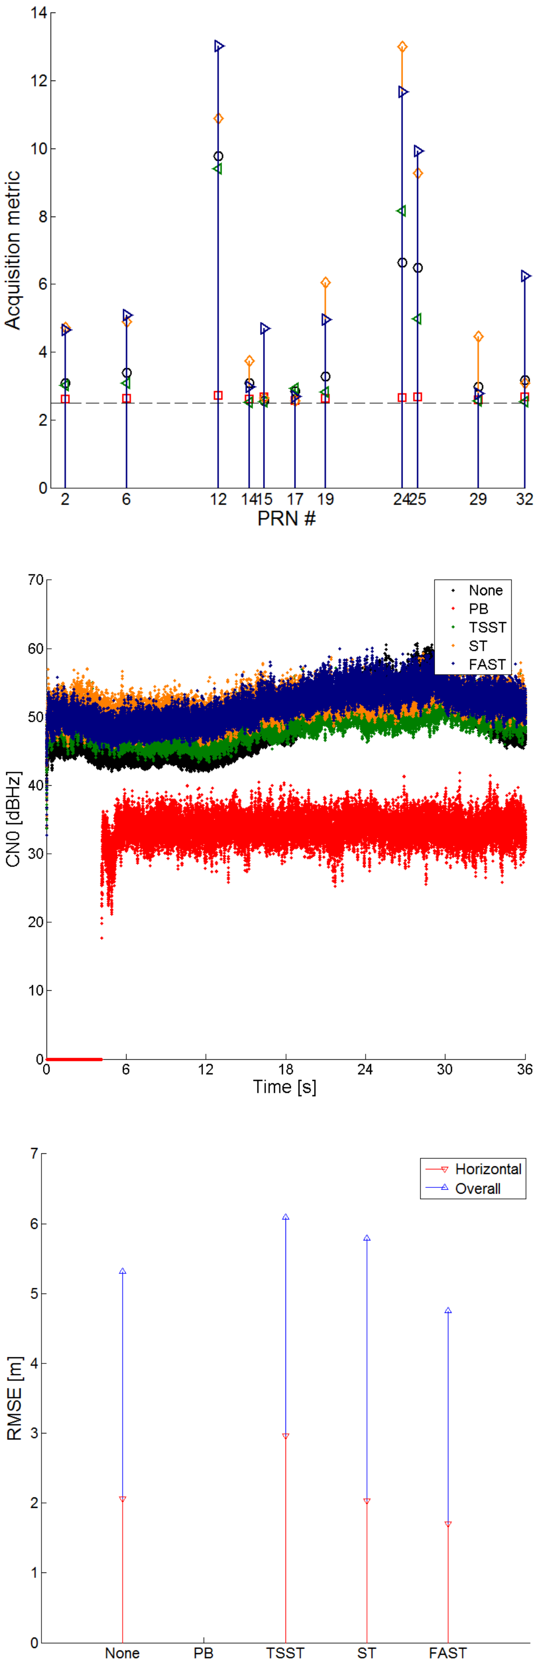

FIGURE 4.14: Acquisition, tracking, and positioning performance in the presence of two jammer with 10dB of JNR overall.

Bibliography

[1] Marco Bartolucci et al. "Joint Jammer Detection and Localization for Dependable GNSS". In: *Proceedings of the ION Pacific PNT Meeting*. 2015, pp. 498–506.

[2] G. Pojani et al. "Optimal EKF for Quasi-Tightly Coupled GNSS/INS Integration". In: *Proceedings of the 9th Annual Baška GNSS Conference*. 2015.

[3] Yazan Adboush et al. "Time-frequency interference rejection based on the S-Transform for GNSS applications". In: *IEEE International Conference on Communications (ICC)*. May 2017.

[4] G. Pojani et al. "Snapshot Localization of a Single Jammer Using TDOA and FDOA Measurements". In: *6th International Colloquium - Scientific and Fundamental Aspects of GNSS / Galileo*. 2017.

[5] M. Bartolucci, G. Pojani, and G. E. Corazza. "Joint jammer detection and localization". In: (2017 - in preparation).

[6] Y. Adboush, Giacomo Pojani, and G. E. Corazza. "A Scalable S-Transform Architecture Based on a Digital Filter Bank With Flexible Sampling". In: (2017 - submitted).

[7] Giacomo Pojani, Yazan Abdoush, and Giovanni Emanuele Corazza. "Performance Assessment of GNSS Receivers Equipped with Jamming Detection and Mitigation Units Based on the S-Transform". In: 2018 - in preparation.

[8] John A. Volpe National Transportation Systems Center. *Vulnerability assessment of the transportation infrastructure relying on the Global Positioning System*. 2001.

[9] G. X. Gao. "DME/TACAN Interference and its Mitigation in L5/E5 Bands". In: *Proceedings of the 20th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS)*. 2007, pp. 1191–1200.

[10] J. Wang. "Pseudolite Applications in Positioning and Navigation: Progress and Problems". In: *Journal of Global Positioning Systems* 1.1 (July 2002), pp. 48–56.

[11] T. E. Humphreys et al. "Assessing the Spoofng Threat: Development of a Portable GPS Civilian Spoofer". In: *Proceedings of the 21st International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS)*. 2008, pp. 2314–2325.

[12] S. Pullen and G. X. Gao. "GNSS jamming in the name of privacy: Potential threat to GPS aviation". In: *Inside GNSS* (2012).

[13] J.C. Grabowski. "Field Observations of Personal Privacy Devices". In: *Proceedings of the International Technical Meeting of The Institute of Navigation (ION GNSS)*. 2012, pp. 689–741.

[14] Thomas Kraus, Roland Bauernfeind, and Bernd Eissfeller. "Survey of In-Car Jammers - Analysis and Modeling of the RF Signals and IF Samples (Suitable for Active Signal Cancelation)". In: *Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2011, pp. 430–435.

[15] G. Gibbons. "FCC Fines Operator of GPS Jammer That Affected Newark Airport GBAS". In: *Inside GNSS* (2013).

[16] K. Sheridan et al. *DETECTOR: Applications and Threats Analysis*. June 2012.

[17] C. Curry. "GPS Jamming: Quantifying the Threat". In: *Workshop on Synchronization in Telecommunication Systems* (2013).

[18] B. Boashash. *Time Frequency Signal Analysis and Processing: A Comprehensive Reference*. Elsevier, 2003.

[19] F. Hlawatsch and G. F. Boudreaux-Bartels. "Linear and quadratic time-frequency signal representations". In: *IEEE Signal Processing Mag.* 9.2 (Apr. 1992), pp. 21–67.

[20] R. A. Brown, M. L. Lauzon, and R. Frayne. "A General Description of Linear Time-Frequency Transforms and Formulation of a Fast, Invertible Transform That Samples the Continuous S-Transform Spectrum Nonredundantly". In: *IEEE Transactions on Signal Processing* 58.1 (Jan. 2010), pp. 281–290.

[21] P. J. Loughlin and L. Cohen. "The uncertainty principle: global, local, or both?" In: *IEEE Transactions on Signal Processing* 52.5 (May 2004), pp. 1218–1227.

[22] J. Allen. "Short term spectral analysis, synthesis, and modification by discrete Fourier transform". In: *IEEE Transactions Acoustics, Speech, Signal Processing* 25.3 (1977), pp. 235–238.

[23] R. G. Stockwell, L. Mansinha, and R. P. Lowe. "Localization of the complex spectrum: the S transform". In: *IEEE Transactions on Signal Processing* 44.4 (Apr. 1996), pp. 998–1001.

[24] F. J. Harris. "On the use of windows for harmonic analysis with the discrete Fourier transform". In: *Proceedings IEEE* 66.1 (Jan. 1978), pp. 51–83.

[25] I. Daubechies. "The wavelet transform, time-frequency localization and signal analysis". In: *IEEE Transactions on Inforamtion Theory* 36.5 (Sept. 1990), pp. 961–1005.

[26] T. Claasen and W. F. G. Mecklenbrauker. "The Wigner distribution-a tool for time-frequency signal analysis. Part II: discrete-time signals". In: *Philips J. Research* 35 (1980), pp. 276–350.

[27] L. Cohen. "Time-frequency distributions-a review". In: *Proceedings IEEE* 77.7 (July 1989), pp. 941–981.

[28] F Beauville et al. "A first comparison of search methods for gravitational wave bursts using LIGO and Virgo simulated data". In: *Classical and Quantum Gravity* 22.18 (2005).

[29] S. Mishra, C. N. Bhende, and B. K. Panigrahi. "Detection and Classification of Power Quality Disturbances Using S-Transform and Probabilistic Neural Network". In: *IEEE Transactions on Power Delivery* 23.1 (Dec. 2008), pp. 280–287.

[30] S. Assous et al. "S-transform applied to laser Doppler flowmetry reactive hyperemia signals". In: *IEEE Transactions on Biomedical Engineering* 53.6 (June 2006), pp. 1032–1037.

[31] A. S. Gonzalez et al. "Non-stationary distributed source approximation: An alternative to improve localization procedures". In: *Human brain mapping* 14.2 (2001), pp. 81–95.

[32] R. G. Stockwell. "A basis for efficient representation of the S-Transform". In: *Digital Signal Processing* 17.1 (2007).

[33] J. B. Allen and L. R. Rabiner. "A unified approach to short-time Fourier analysis and synthesis". In: *Proceedings IEEE* 65.11 (Nov. 1977), pp. 1558–1564.

[34] M. Biswal and P. K. Dash. "Detection and characterization of multiple power quality disturbances with a fast S-transform and decision tree based classifier". In: *Digital Signal Processing* 23.4 (June 2013), pp. 1071–1083.

[35] M. Schimmel and J. Gallart. "The inverse S-transform in filters with time-frequency localization". In: *IEEE Transactions on Signal Processing* 53.11 (Nov. 2005), pp. 4417–4422.

[36] C. Simon et al. "The S-Transform and Its Inverses: Side Effects of Discretizing and Filtering". In: *IEEE Transactions on Signal Processing* 55.10 (Oct. 2007), pp. 4928–4937.

[37] S. C. Pei and P. W. Wang. "Novel Inverse S transform With Equalization Filter". In: *IEEE Transactions on Signal Process* 57.10 (Oct. 2009), pp. 3858–3868.

[38] S. Hamid Nawab and Thomas F. Quatieri. "Advanced Topics in Signal Processing". In: ed. by J. S. Lim and Alan V. Oppenheim. Prentice-Hall, Inc., 1988. Chap. Short-Time Fourier Transform, pp. 289–338.

[39] W. G. Gardner. "Efficient Convolution without Input-Output Delay". In: *Journal of the Audio Engineering Society* 43.3 (Mar. 1995), pp. 127–136.

[40] C. R. Pinnegar and L. Mansinha. "The S-transform with windows of arbitrary and varying shape". In: *Geophysics* 68.1 (Jan. 2003), pp. 381–385.

[41] S. C. Pei and S. G. Huang. "STFT With Adaptive Window Width Based on the Chirp Rate". In: *IEEE Transactions on Signal Processing* 60.8 (Aug. 2012), pp. 4065–4080.

[42] V. Katkovnik and L. J. Stankovic. "Periodogram with varying and data-driven window length". In: *Signal Processing* 67.3 (June 1998), pp. 345–358.

[43] D. Borio et al. "Time-Frequency Excision for GNSS Applications". In: *IEEE Systems Journal* 2.1 (Mar. 2008), pp. 27–37.

[44] M. Portnoff. "Implementation of the digital phase vocoder using the fast Fourier transform". In: *IEEE Transactions Acoustics, Speech, Signal Processing* 24.3 (June 1976), pp. 243–248.

[45] B. Boashash. "Estimating and interpreting the instantaneous frequency of a signal. II. Algorithms and applications". In: *Proceedings IEEE* 80.4 (Apr. 1992), pp. 540–568.

[46] V. N. Ivanovic, M. Dakovic, and L. Stankovic. "Performance of quadratic time-frequency distributions as instantaneous frequency estimators". In: *IEEE Transactions on Signal Processing* 51.1 (2003), pp. 77–89.

[47] K. M. Wong and Q. Jin. "Estimation of the time-varying frequency of a signal: the Cramer-Rao bound and the application of Wigner distribution". In: *IEEE Transactions Acoustics, Speech, Signal Processing* 38.3 (Mar. 1990), pp. 519–536.

[48] V. Katkovnik and L. Stankovic. “Instantaneous frequency estimation using the Wigner distribution with varying and data-driven window length”. In: *IEEE Transactions on Signal Processing* 46.9 (Sept. 1998), pp. 2315–2325.

[49] A. J. Weiss. “Direct Geolocation of Wideband Emitters Based on Delay and Doppler”. In: *IEEE Transactions on Signal Processing* 59.6 (June 2011), pp. 2513–2521.

[50] D. J. Torrieri. “Statistical Theory of Passive Location Systems”. In: *IEEE Transactions on Aerospace and Electronic Systems* 20.2 (Mar. 1984), pp. 183–198.

[51] Y. T. Chan and K. C. Ho. “A simple and efficient estimator for hyperbolic location”. In: *IEEE Transactions on Signal Processing* 42.8 (Aug. 1994), pp. 1905–1915.

[52] W. H. Foy. “Position-Location Solutions by Taylor-Series Estimation”. In: *IEEE Transactions on Aerospace and Electronic Systems* 12.2 (Mar. 1976), pp. 187–194.

[53] A. Amar and A. J. Weiss. “Localization of Narrowband Radio Emitters Based on Doppler Frequency Shifts”. In: *IEEE Transactions on Signal Processing* 56.11 (Nov. 2008), pp. 5500–5508.

[54] Y. T. Chan and J. J. Towers. “Sequential localization of a radiating source by Doppler-shifted frequency measurements”. In: *IEEE Transactions on Aerospace and Electronic System* 28.4 (Oct. 1992), pp. 1084–1090.

[55] K. C. Ho and Wenwei Xu. “An accurate algebraic solution for moving source location using TDOA and FDOA measurements”. In: *IEEE Transactions on Signal Processing* 52.9 (Sept. 2004), pp. 2453–2463.

[56] K. C. Ho, X. Lu, and L. Kovavisaruch. “Source Localization Using TDOA and FDOA Measurements in the Presence of Receiver Location Errors: Analysis and Solution”. In: *IEEE Transactions on Signal Processing* 55.2 (Feb. 2007), pp. 684–696.

[57] A. Beck, P. Stoica, and J. Li. “Exact and Approximate Solutions of Source Localization Problems”. In: *IEEE Transactions on Signal Processing* 56.5 (May 2008), pp. 1770–1778.

[58] G. Wang et al. “A Bias-Reduced Nonlinear WLS Method for TDOA/FDOA-Based Source Localization”. In: *IEEE Transactions on Vehicular Technology* 65.10 (Oct. 2016), pp. 8603–8615.

[59] X. Qu, L. Xie, and W. Tan. “Iterative Constrained Weighted Least Squares Source Localization Using TDOA and FDOA Measurements”. In: *IEEE Transactions on Signal Processing* 65.15 (Aug. 2017), pp. 3990–4003.

[60] G. Wang and H. Chen. “An Importance Sampling Method for TDOA-Based Source Localization”. In: *IEEE Transactions on Wireless Communications* 10.5 (May 2011), pp. 1560–1568.

[61] G. Wang et al. “A Semidefinite Relaxation Method for Source Localization Using TDOA and FDOA Measurements”. In: *IEEE Transactions on Vehicular Technology* 62.2 (Feb. 2013), pp. 853–862.

[62] H. W. Wei et al. “Multidimensional Scaling Analysis for Passive Moving Target Localization With TDOA and FDOA Measurements”. In: *IEEE Transactions on Signal Processing* 58.3 (Mar. 2010), pp. 1677–1688.

[63] N. Okello et al. “Comparison of Recursive Algorithms for Emitter Localisation using TDOA Measurements from a Pair of UAVs”. In: *IEEE Transactions on Aerospace and Electronic Systems* 47.3 (Jan. 2011), pp. 1723–1732.

[64] D. Musicki, R. Kaune, and W. Koch. “Mobile Emitter Geolocation and Tracking Using TDOA and FDOA Measurements”. In: *IEEE Transactions on Signal Processing* 58.3 (Mar. 2012), pp. 1863–1874.

[65] H. W. Sorenson. “Least-squares estimation: from Gauss to Kalman”. In: *IEEE Spectrum* 7.7 (July 1970), pp. 63–68.

[66] Li Cong and Weihua Zhuang. “Nonline-of-sight error mitigation in mobile location”. In: *IEEE Transactions on Wireless Communications* 4.2 (Mar. 2005), pp. 560–573.

[67] Yiu-Tong Chan et al. “Time-of-arrival based localization under NLOS conditions”. In: *IEEE Transactions on Vehicular Technology* 55.1 (Jan. 2006), pp. 17–24.

[68] W. Xu et al. “Distributed Localization of a RF Target in NLOS Environments”. In: *IEEE Journal on Selected Areas in Communications* 33.7 (July 2015), pp. 1317–1330.

[69] R. J. R. Thompson, E. Cetin, and A. G. Dempster. “Unknown source localization using RSS in open areas in the presence of ground reflections”. In: *Proceedings of the IEEE/ION Position, Location and Navigation Symposium (PLANS)*. 2012, pp. 1018–1027.

[70] B. R. Jackson, S. Wang, and R. Inkol. “Received signal strength difference emitter geolocation least squares algorithm comparison”. In: *Canadian Conference on Electrical and Computer Engineering (CCECE)*. 2011.

[71] K. Spingarn. “Passive Position Location Estimation Using the Extended Kalman Filter”. In: *IEEE Transactions on Aerospace and Electronic Systems* 23.4 (July 1987), pp. 558–567.

[72] M. Gavish and A. J. Weiss. “Performance analysis of bearing-only target location algorithms”. In: *IEEE Transactions on Aerospace and Electronic Systems* 28.3 (July 1992), pp. 817–828.

[73] K. Becker. “An efficient method of passive emitter location”. In: *IEEE Transactions on Aerospace and Electronic Systems* 28.4 (Oct. 1992), pp. 1091–1104.

[74] K. Becker. “Three-dimensional target motion analysis using angle and frequency measurements”. In: *IEEE Transactions on Aerospace and Electronic Systems* 41.1 (2005), pp. 284–301.

[75] Li Cong and Weihua Zhuang. “Hybrid TDOA/AOA mobile user location for wideband CDMA cellular systems”. In: *IEEE Transactions on Wireless Communications* 1.3 (July 2002), pp. 439–447.

[76] Changlin Ma, R. Klukas, and G. Lachapelle. “An enhanced two-step least squared approach for TDOA/AOA wireless location”. In: *IEEE International Conference on Communications (ICC)*. 2003.

[77] S. D. Coutts. “Passive localization of moving emitters using out-of-plane multipath”. In: *IEEE Transactions on Aerospace and Electronic Systems* 36.2 (Apr. 2000), pp. 584–595.

[78] A. O’connor, P. Setlur, and N. Devroye. “Single-sensor RF emitter localization based on multipath exploitation”. In: *IEEE Transactions on Aerospace and Electronic Systems* 51.3 (July 2015), pp. 1635–1651.

[79] R.O. Schmidth. “Muliple Emitter Location and Signal Parameter Estimation”. In: *IEEE Transacations on Antennas and Propagation* 34.3 (Mar. 1986), pp. 276–280.

[80] A. Hu et al. “An ESPRIT-Based Approach for 2-D Localization of Incoherently Distributed Sources in Massive MIMO Systems”. In: *IEEE Journal of Selected Topics in Signal Processing* 8.5 (Oct. 2014), pp. 996–1011.

[81] W. T. Zhang et al. “Tracking Multiple Targets in MIMO Radar Via Adaptive Asymmetric Joint Diagonalization”. In: *IEEE Transactions on Signal Processing* 64.11 (June 2016), pp. 2880–2893.

[82] S. Qin, Y. D. Zhang, and M. G. Amin. “Generalized Coprime Array Configurations for Direction-of-Arrival Estimation”. In: *IEEE Transactions on Signal Processing* 63.6 (Mar. 2015), pp. 1377–1390.

[83] T. Sathyan, A. Sinha, and T. Kirubarajan. “Passive geolocation and tracking of an unknown number of emitters”. In: *IEEE Transactions on Aerospace and Electronic Systems* 42.2 (Apr. 2006), pp. 740–750.

[84] D. Musicki. “Multi-target tracking using multiple passive bearings-only asynchronous sensors”. In: *IEEE Transactions on Aerospace and Electronic Systems* 44.3 (July 2008), pp. 1151–1160.

[85] D. Carevic. “Automatic Estimation of Multiple Target Positions and Velocities Using Passive TDOA Measurements of Transients”. In: *IEEE Transactions on Signal Processing* 55.2 (Feb. 2007), pp. 424–436.

[86] M. Sun and K. C. Ho. “An Asymptotically Efficient Estimator for TDOA and FDOA Positioning of Multiple Disjoint Sources in the Presence of Sensor Location Uncertainties”. In: *IEEE Transactions on Signal Processing* 59.7 (July 2011), pp. 3434–3440.

[87] C. Hue, J. P. Le Cadre, and P. Perez. “Sequential Monte Carlo methods for multiple target tracking and data fusion”. In: *IEEE Transactions on Signal Processing* 50.2 (Feb. 2002), pp. 309–325.

[88] Pau Closas. “Bayesian Signal Processing Techniques for GNSS Receivers: from multipath mitigation to positioning”. PhD thesis. Universitat Politecnica de Catalunya, 2009.

[89] Z. Chen. “Bayesian filtering: From Kalman filters to particle filters, and beyond”. In: *Statistics* 182.1 (2003), pp. 1–69.

[90] L. Jiang, S. S. Singh, and S. Yildirim. “A new particle filtering algorithm for multiple target tracking with non-linear observations”. In: *International Conference on Information Fusion (FUSION)*. 2014.

[91] K. Yang, G. Wang, and Z. Q. Luo. “Efficient Convex Relaxation Methods for Robust Target Localization by a Sensor Network Using Time Differences of Arrivals”. In: *IEEE Transactions on Signal Processing* 57.7 (July 2009), pp. 2775–2784.

[92] O. Jean and A. J. Weiss. “Passive Localization and Synchronization Using Arbitrary Signals”. In: *IEEE Transactions on Signal Processing* 62.8 (Apr. 2014), pp. 2143–2150.

[93] Y. Zhang et al. “Distributed Projection-Based Algorithms for Source Localization in Wireless Sensor Networks”. In: *IEEE Transactions on Wireless Communications* 14.6 (June 2015), pp. 3131–3142.

[94] Y. I. Wu, H. Wang, and X. Zheng. “WSN Localization Using RSS in Three-Dimensional Space—A Geometric Method With Closed-Form Solution”. In: *IEEE Sensors Journal* 16.11 (June 2016), pp. 4397–4404.

[95] T. L. T. Nguyen et al. "A Bayesian Perspective on Multiple Source Localization in Wireless Sensor Networks". In: *IEEE Transactions on Signal Processing* 64.7 (Apr. 2016), pp. 1684–1699.

[96] N. Patwari et al. "Relative location estimation in wireless sensor networks". In: *IEEE Transactions on Signal Processing* 51.8 (Aug. 2003), pp. 2137–2148.

[97] S. Hara and D. Anzai. "Experimental Performance Comparison of RSSI- and TDOA-Based Location Estimation Methods". In: *IEEE Vehicular Technology Conference (VTC)*. 2008.

[98] A. Mpitziopoulos et al. "A survey on jamming attacks and countermeasures in WSNs". In: *IEEE Communications Surveys & Tutorials* 11.4 (2009), pp. 42–56.

[99] K. Grover, A. Lim, and Q. Yang. "Jamming and anti-jamming techniques in wireless networks: a survey". In: *International Journal of Ad Hoc and Ubiquitous Computing* 17.4 (Dec. 2014), pp. 197–215.

[100] K. Pelechrinis et al. "Lightweight Jammer Localization in Wireless Networks: System Design and Implementation". In: *IEEE Global Telecommunications Conference (GLOBECOM)*. 2009.

[101] Z. Liu et al. "Exploiting Jamming-Caused Neighbor Changes for Jammer Localization". In: *IEEE Transactions on Parallel and Distributed Systems* 23.3 (Mar. 2012), pp. 547–555.

[102] Hongbo Liu et al. "Localizing jammers in wireless networks". In: *IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 2009.

[103] T. Cheng, P. Li, and S. Zhu. "Multi-jammer Localization in Wireless Sensor Networks". In: *International Conference on Computational Intelligence and Security (CIS)*. 2011.

[104] H. Liu et al. "Localizing Multiple Jamming Attackers in Wireless Networks". In: *International Conference on Distributed Computing Systems*. 2011.

[105] T. Cheng, P. Li, and S. Zhu. "An Algorithm for Jammer Localization in Wireless Sensor Networks". In: *IEEE International Conference on Advanced Information Networking and Applications (AINA)*. 2012.

[106] Z. Liu et al. "An Error-Minimizing Framework for Localizing Jammers in Wireless Networks". In: *IEEE Transactions on Parallel and Distributed Systems* 25.2 (Feb. 2014), pp. 508–517.

[107] M. Bshara et al. "Fingerprinting Localization in Wireless Networks Based on Received-Signal-Strength Measurements: A Case Study on WiMAX Networks". In: *IEEE Transactions on Vehicular Technology* 59.1 (Jan. 2010), pp. 283–294.

[108] S. Tomic, M. Beko, and R. Dinis. "RSS-Based Localization in Wireless Sensor Networks Using Convex Relaxation: Noncooperative and Cooperative Schemes". In: *IEEE Transactions on Vehicular Technology* 64.5 (May 2015), pp. 2037–2050.

[109] A. Haniz, G. K. Tran, and K. Saito. "A Novel Phase-Difference Fingerprinting Technique for Localization of Unknown Emitters". In: *IEEE Transactions on Vehicular Technology* 66.9 (Sept. 2017), pp. 8445–8457.

[110] A. G. Dempster and E. Cetin. "Interference Localization for Satellite Navigation Systems". In: *Proceedings of the IEEE* 104.6 (June 2016), pp. 1318–1326.

[111] Alison Brown et al. “Jammer and Interference Location”. In: *Proceedings of the 12th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS)*. 1999, pp. 137–142.

[112] M. Trinkle and D.A. Gray. “Interference Localisation Trials Using Adaptive Antenna Arrays”. In: *Proceedings of the 15th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS)*. 2002, pp. 613–619.

[113] Jonas Lindstrom et al. “GNSS Interference Detection and Localization using a Network of Low Cost Front-End Modules”. In: *Proceedings of the 20th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS)*. 2007, pp. 1165–1172.

[114] Logan Scott. “J911: The Case for Fast Jammer Detection and Location Using Crowdsourcing Approaches”. In: *Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2011, pp. 1931–1940.

[115] R. Bauernfeind et al. “In-Car Jammer interference detection in automotive GNSS receivers and localization by means of vehicular communication”. In: *IEEE Forum on Integrated and Sustainable Transportation Systems*. 2011.

[116] Daniele Borio et al. “Jammer Localization: From Crowdsourcing to Synthetic Detection”. In: *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+)*. 2016, pp. 3107–3116.

[117] R. Yozevitch, B. Ben Moshe, and S. Safrigin. “Tackling the GNSS jamming problem using a particle filter algorithm”. In: *IEEE 28th Convention of Electrical & Electronics Engineers in Israel (IEEEI)*. 2014.

[118] R. J. R. Thompson, A. T. Balaei, and A. G. Dempster. “Outdoor localization of a WiFi source with unknown transmission power”. In: *International Global Navigation Satellite Systems Society Symposium (IGNSS)*. 2009.

[119] M. Bartolucci et al. “Cooperative/distributed localization and characterization of GNSS jamming interference”. In: *International Conference on Localization and GNSS (ICL-GNSS)*. 2013.

[120] E. Cetin, R. J. R. Thompson, and A. G. Dempster. “Interference Localisation within the GNSS Environmental Monitoring System (GEMS)”. In: *International Global Navigation Satellite Systems Society Symposium (IGNSS)*. 2011.

[121] Matthew Trinkle et al. “Interference Localisation within the GNSS Environmental Monitoring System (GEMS) - Initial Field Test Results”. In: *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2012, pp. 2930–2939.

[122] E. Cetin, R. J. R. Thompson, and A. G. Dempster. “Passive interference localization within the GNSS environmental monitoring system (GEMS): TDOA aspects”. In: *GPS Solutions* 18.4 (2014), pp. 483–495. ISSN: 1521-1886.

[123] Adrien Perkins et al. “Demonstration of UAV Based GPS Jammer Localization During a Live Interference Exercise”. In: *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+)*. 2016, pp. 3094–3106.

[124] Zhonghai Wang et al. “Jamming emitter localization with multiple UAVs equipped with smart antennas”. In: *Proceedings of SPIE*. Vol. 7696. 2010.

[125] F. D. Nunes and F. M. G. Sousa. "GNSS Near-Far Mitigation through Subspace Projection without Phase Information". In: *IEEE Transactions on Aerospace and Electronic Systems* 48.3 (July 2012), pp. 2746–2755.

[126] Z. Xu, M. Trinkle, and D. A. Gray. "Weak interference direction of arrival estimation in the GPS L1 frequency band". In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2015.

[127] A. Bours, E. Cetin, and A. G. Dempster. "Enhanced GPS interference detection and localisation". In: *Electronics Letters* 50.19 (Sept. 2014), pp. 1391–1393.

[128] G. Gabelli et al. "GNSS Signal Cancellation for Enhanced Interference Detection and Localization". In: *International Global Navigation Satellite Systems Society Symposium (IGNSS)*. 2013.

[129] R. J. R. Thompson, E. Cetin, and A. G. Dempster. "Influence of GPS Satellites Cross-Correlation on the TDOA Measurements within the GNSS Environmental Monitoring System (GEMS)". In: *International Global Navigation Satellite Systems Society Symposium (IGNSS)*. 2011.

[130] Ryan J.R Thompson, Ediz Cetin, and Andrew G. Dempster. "Evaluation of Relative GPS Timing Under Jamming Conditions". In: *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2012, pp. 717–730.

[131] R. H. Mitch et al. "Civilian GPS jammer signal tracking and geolocation". In: *Proceedings 25th International Techical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2012, pp. 2901–2920.

[132] Oscar Isoz, Asghar T. Balaei, and Dennis Akos. "Interference Detection and Localization in the GPS L1 Band". In: *Proceedings of the International Technical Meeting of The Institute of Navigation (ION GNSS)*. 2010, pp. 925–929.

[133] O. Isoz and D. Akos. "Development of a deployable low cost interference detection and localization system for the GNSS L1/E1 band". In: *5th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*. 2010.

[134] J. P. Poncelet and D. M. Akos. "A low-cost monitoring station for detection & localization of interference in GPS L1 band". In: *6th ESA Workshop on Satellite Navigation Technologies (NAVITEC) & European Workshop on GNSS Signals and Signal Processing*. 2012.

[135] Konstantin Gromov et al. "GIDL: Generalized Interference Detection and Localization System". In: *Proceedings of the 13th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS)*. 2000, pp. 447–457.

[136] J. A. Bhatti, T. E. Humphreys, and B. M. Ledvina. "Development and demonstration of a TDOA-based GNSS interference signal localization system". In: *Proceedings of the IEEE/ION Position, Location and Navigation Symposium*. 2012, pp. 455–469.

[137] W. U. Bajwa, K. Gedalyahu, and Y. C. Eldar. "Identification of Parametric Underspread Linear Systems and Super-Resolution Radar". In: *IEEE Transactions on Signal Processing* 59.6 (June 2011), pp. 2548–2561.

[138] Kenneth M. Pesyna et al. "Tightly-Coupled Opportunistic Navigation for Deep Urban and Indoor Positioning". In: *Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2011, pp. 3605–3616.

[139] R. J. R. Thompson, A. T. Balaei, and A. G. Dempster. "Dilution of precision for GNSS interference localisation systems". In: *European Navigation Conference (ENC GNSS)*. 2009.

[140] J. A. Garcia-Molina and M. Crisci. "Snapshot localisation of multiple jammers based on receivers of opportunity". In: *8th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*. 2016.

[141] M. Bartolucci et al. "Synchronisation of low-cost open source SDRs for navigation applications". In: *8th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*. 2016.

[142] J. A. Garcia-Molina and M. Crisci. "Cloud-based Localization of GNSS Jammers". In: *Proceedings of the 28th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+)*. 2015, pp. 3289–3295.

[143] Oscar Isoz et al. "Assessment of GPS L1/Galileo E1 Interference Monitoring System for the Airport Environment". In: *Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2011, pp. 1920–1930.

[144] M. Cuntz et al. "GALANT - Galileo Antenna and Receiver Demonstrator for Safety-Critical Applications". In: *European Conference on Wireless Technologies*. 2007, pp. 59–61.

[145] M. Cuntz et al. "A Multi Antenna Receiver for Galileo SoL Applications". In: *IEEE International Mini-Symposium on Electromagnetics and Network Theory and their Microwave Technology Applications (MTT-S)*. 2008.

[146] R. H. Mitch et al. "Signal Characteristics of Civil GPS Jammers". In: *Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2011, pp. 1907–1919.

[147] S. Stein. "Algorithms for ambiguity function processing". In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 29.3 (June 1981), pp. 588–599.

[148] A. Yeredor and E. Angel. "Joint TDOA and FDOA Estimation: A Conditional Bound and Its Use for Optimally Weighted Localization". In: *IEEE Transactions on Signal Processing* 59.4 (Apr. 2011), pp. 1612–1623.

[149] S. J. Julier and J. K. Uhlmann. "Unscented Filtering and Nonlinear Estimation". In: *Proceedings of the IEEE* 92.3 (Mar. 2004), pp. 401–422.

[150] M. Roth and F. Gustafsson. "An efficient implementation of the second order extended Kalman filter". In: *14th International Conference on Information Fusion (FUSION)*. 2011.

[151] F. Gustafsson and G. Hendeby. "Some Relations Between Extended and Unscented Kalman Filters". In: *IEEE Transacations on Signal Processing* 60.2 (Feb. 2012), pp. 545–547.

[152] T. Kirubarajan and Y. Bar-Shalom. "Kalman filter versus IMM estimator: when do we need the latter?" In: *IEEE Transactions on Aerospace and Electronic Systems* 39.4 (Oct. 2003), pp. 1452–1457.

[153] P. Tichavsky, C. H. Muravchik, and A. Nehorai. "Posterior Cramer-Rao bounds for discrete-time nonlinear filtering". In: *IEEE Transacations on Signal Processing* 46.5 (May 1998), pp. 1386–1396.

[154] C. Fritsche et al. “A fresh look at Bayesian Cramer-Rao bounds for discrete-time nonlinear filtering”. In: *17th International Conference on Information Fusion (FUSION)*. 2014.

[155] R. Van der Merwe and E. A. Wan. “The square-root unscented Kalman filter for state and parameter-estimation”. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. Vol. 6. 2001, pp. 3461–3464.

[156] S. S. Sahu, G. Panda, and N. V. George. “An Improved S-Transform for Time-Frequency Analysis”. In: *IEEE International Advance Computing Conference*. 2009.

[157] S. C. Pei and P. W. Wang. “Energy concentration enhancement using window width optimization in S transform”. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*. 2010.

[158] S. R. Deans. “Hough Transform from the Radon Transform”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 3.2 (Mar. 1981), pp. 185–188.

[159] Michail Vlachos, Philip Yu, and Vittorio Castelli. “On Periodicity Detection and Structural Periodic Similarity”. In: *Proceedings of the 2005 SIAM International Conference on Data Mining*. 2005, pp. 449–460.

[160] M. Vlachos et al. “Identifying similarities, periodicities and bursts for online search queries”. In: *Proceedings of the 2004 ACM SIGMOD international conference on management of data*. 2004, pp. 131–142.

[161] G. Peeters. “Music Pitch Representation by Periodicity Measures Based on Combined Temporal and Spectral Representations”. In: *IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*. 2006.

[162] G. Peeters. “Spectral and Temporal Periodicity Representations of Rhythm for the Automatic Classification of Music Audio Signal”. In: *IEEE Transactions on Audio, Speech, and Language Processing* 19.5 (July 2011), pp. 1242–1252.

[163] M. Ester et al. “A density-based algorithm for discovering clusters a density-based algorithm for discovering clusters in large spatial databases with noise”. In: *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD)*. 1996, pp. 226–231.

[164] G. Pojani et al. “Characterization and Localization of Multiple In-Car Jammers Based on Time-Frequency Analysis”. In: (2018 - in preparation).

[165] D. Borio, C. O’Driscoll, and J. Fortuny. “Jammer impact on Galileo and GPS receivers”. In: *International Conference on Localization and GNSS (ICL-GNSS)*. 2013.

[166] A. T. Balaei, A. G. Dempster, and D. Akos. “Quantization Degradation of GNSS Signal Quality in the Presence of CW RFI”. In: *IEEE 10th International Symposium on Spread Spectrum Techniques and Applications*. 2008.

[167] F. Bastide et al. “Automatic gain control (AGC) as an interference assessment tool”. In: *Proceedings of the 16th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2003, 2042–2053.

[168] A. Balaei. “Statistical inference technique in pre-correlation interference detection in GPS receivers”. In: *Proceedings of the 19th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*. 2006, 2232–2240.

[169] A. Tani and R. Fantacci. “Performance evaluation of a precorrelation interference detection algorithm for the GNSS based on nonparametrical spectral estimation”. In: *IEEE Systems Journal* 2.1 (Mar. 2008), 20–26.

[170] R. Calcagno et al. “An interference detection algorithm for COTS GNSS receivers”. In: *5th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*. 2010.

[171] B. Motella et al. “Method for Assessing the Interference Impact on GNSS Receivers”. In: *IEEE Transactions on Aerospace and Electronic Systems* 47.2 (Apr. 2011), pp. 1416–1432.

[172] F. D. Nunes and F. M. G. Sousa. “Jamming detection in GNSS signals using the sample covariance matrix”. In: *6th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*. 2012.

[173] B. Motella, M. Pini, and L. L. Presti. “GNSS interference detector based on chi-square goodness-of-fit test”. In: *6th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*. 2012.

[174] E. Axell et al. “Jamming detection in GNSS receivers: Performance evaluation of field trials”. In: *Proceedings of the 26th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+)*. 2013, pp. 2542–2551.

[175] D. Borio and C. Gioia. “Real-time jamming detection using the sum-ofsquares paradigm”. In: *International Conference on Localization and GNSS (ICL-GNSS)*. 2015.

[176] K. Sun, M. Zhang, and D. Yang. “A new interference detection method based on hybrid time-frequency distribution for GNSS receivers”. In: *IEEE Transactions on Vehicular Technology* 65.11 (Nov. 2016), pp. 9057–9071.

[177] H. Kuusniemi. “Effects of GNSS jammers and potential mitigation approaches”. In: *United Nations and Latvia Workshop on the Applications of GNSS*. 2012.

[178] Grace Xingxin Gao et al. “Protecting GNSS receivers from jamming and interference”. In: *Proceedings of the IEEE* 104.6 (2016), pp. 1327–1338.

[179] M. G. Amin. “Interference mitigation in spread spectrum communication systems using time-frequency distributions”. In: *IEEE Transactions on Signal Processing* 45.1 (Jan. 1997), pp. 90–101.

[180] S. Barbarossa and A. Scaglione. “Adaptive time-varying cancellation of wideband interferences in spread-spectrum communications based on time-frequency distributions”. In: *IEEE Transactions on Signal Processing* 47.4 (Apr. 1999), pp. 957–965.

[181] Daniele Borio, Laura Camoriano, and Letizia Lo Presti. “Two-pole and multipole notch filters: a computationally effective solution for GNSS interference detection and mitigation”. In: *IEEE Systems Journal* 2.1 (2008), pp. 38–47.

[182] D. Borio, C. O’Driscoll, and J. Fortuny. “GNSS Jammers: Effects and Countermeasures”. In: *6th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*. 2012.

[183] D. Borio, J. Fortuny, and C. O’Driscoll. “Spectral and spatial characterization of GNSS jammers”. In: *GNSS Vulnerabilities Solutions Conference*. 2013.

[184] Y. Chien. "Design of GPS anti-jamming systems using adaptive notch filters". In: *IEEE Systems Journal* 9.2 (Oct. 2015), pp. 451–460.

[185] Daniele Borio. "Swept GNSS jamming mitigation through pulse blanking". In: *European Navigation Conference (ENC)*. 2016.

[186] P. T. Capozza et al. "A single-chip narrow-band frequency-domain excisor for a global positioning system (GPS) receiver". In: *IEEE Journal of Solid-State Circuits* 35.3 (Mar. 2000), pp. 401–411.

[187] Y. Zhang, M. G. Amin, and A. R. Lindsey. "Mitigation of periodic interferers in GPS receivers using subspace projection techniques". In: *Proceedings of the Sixth International Symposium on Signal Processing and its Applications*. 2001, pp. 497–500.

[188] Y. Zhang, M. G. Amin, and A. R. Lindsey. "Anti-jamming GPS receivers based on bilinear signal distributions". In: *IEEE Military Communications Conference (MILCOM)*. 2001, pp. 1070–1074.

[189] Simone Savasta, Letizia Lo Presti, and Marco Rao. "Interference mitigation in GNSS receivers by a time-frequency approach". In: *IEEE Transactions on Aerospace and Electronic Systems* 49.1 (Jan. 2013), pp. 415–438.

[190] X. Ouyang and M. G. Amin. "Short-time Fourier transform receiver for non-stationary interference excision in direct sequence spread spectrum communications". In: *IEEE Transactions on Signal Processing* 49.4 (Apr. 2001), pp. 851–863.

[191] S. Barbarossa and O. Lemoine. "Analysis of nonlinear FM signals by pattern recognition of their time-frequency representation". In: *IEEE Signal Processing Letters* 3.4 (Apr. 1996), pp. 112–115.

[192] R. Bauernfeind et al. *Analysis, Detection and Mitigation of In-car GNSS jammer Interference in Intelligent Transport Systems*. Deutscher Luft-und Raumfahrtkongress (DLR). 2012.

[193] K. Borre et al. *A Software-Defined GPS and Galileo Receiver: A single-frequency approach*. Springer, 2007.

[194] D. Borio and L. Lo Presti. "Data and Pilot Combining for Composite GNSS Signal Acquisition". In: *International Journal of Navigation and Observation* (2008).

[195] C. Fernandez-Prades et al. "An open source Galileo E1 software receiver". In: *6th ESA Workshop on Satellite Navigation Technologies (NAVITEC)*. 2012.