

Efficient Schemes for Adaptive Frequency Tracking and their Relevance for EEG and ECG

THÈSE N° 5476 (2012)

PRÉSENTÉE LE 12 OCTOBRE 2012

À LA FACULTÉ DES SCIENCES ET TECHNIQUES DE L'INGÉNIEUR
LABORATOIRE DE TRAITEMENT DES SIGNAUX 1
PROGRAMME DOCTORAL EN GÉNIE ÉLECTRIQUE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Jérôme VAN ZAEN

acceptée sur proposition du jury:

Prof. P. Vandergheynst, président du jury
Dr J.-M. Vesin, directeur de thèse
Dr P. Gonçalves, rapporteur
Dr H. Lissek, rapporteur
Prof. R. Reilly, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2012

*« If you try and take a cat apart to see
how it works, the first thing you have on
your hands is a non-working cat. »*

Douglas Adams

Abstract

Amplitude and frequency are the two primary features of one-dimensional signals, and thus both are widely utilized to analysis data in numerous fields. While amplitude can be examined directly, frequency requires more elaborate approaches, except in the simplest cases. Consequently, a large number of techniques have been proposed over the years to retrieve information about frequency. The most famous method is probably power spectral density estimation. However, this approach is limited to stationary signals since the temporal information is lost. Time-frequency approaches were developed to tackle the problem of frequency estimation in non-stationary data. Although they can estimate the power of a signal in a given time interval and in a given frequency band, these tools have two drawbacks that make them less valuable in certain situations. First, due to their interdependent time and frequency resolutions, improving the accuracy in one domain means decreasing it in the other one. Second, it is difficult to use this kind of approach to estimate the instantaneous frequency of a specific oscillatory component. A solution to these two limitations is provided by adaptive frequency tracking algorithms. Typically, these algorithms use a time-varying filter (a band-pass or notch filter in most cases) to extract an oscillation, and an adaptive mechanism to estimate its instantaneous frequency.

The main objective of the first part of the present thesis is to develop such a scheme for adaptive frequency tracking, the single frequency tracker. This algorithm compares favorably with existing methods for frequency tracking in terms of bias, variance and convergence speed. The most distinguishing feature of this adaptive algorithm is that it maximizes the oscillatory behavior at its output. Furthermore, due to its specific time-varying band-pass filter, it does not introduce any distortion in the extracted component. This scheme is also extended to tackle certain situations, namely the presence of several oscillations in a single signal, the related issue of harmonic components, and the availability of more than one signal with the oscillation of interest. The first extension is aimed at tracking several components simultaneously. The basic idea is to use one tracker to estimate the instantaneous frequency of each oscillation. The second extension uses the additional information contained in several signals to achieve better overall performance. Specifically, it computes separately instantaneous frequency estimates for all available signals which are then combined with weights minimizing the estimation variance. The third extension, which is based on an idea similar to the first one and uses the same weighting procedure as the second one, takes into account the harmonic structure of a signal to improve the estimation performance. A non-causal iterative method for offline processing is also developed in order to enhance an initial frequency trajectory by using future information in addition to past information. Like the single frequency tracker, this method aims at maximizing the oscillatory behavior at the output. Any approach can be used to obtain the initial trajectory.

In the second part of this dissertation, the schemes for adaptive frequency tracking developed in the first part are applied to electroencephalographic and electrocardiographic data. In a first study, the single frequency tracker is used to analyze interactions between neuronal oscillations in different frequency bands, known as cross-frequency couplings, during a visual evoked potential experiment with illusory contour stimuli. With this adaptive approach ensuring that meaningful phase information is extracted, the differences in coupling strength between stimuli

with and without illusory contours are more clearly highlighted than with traditional methods based on predefined filter-banks. In addition, the adaptive scheme leads to the detection of differences in instantaneous frequency. In a second study, two organization measures are derived from the harmonic extension. They are based on the power repartition in the frequency domain for the first one and on the phase relation between harmonic components for the second one. These measures, computed from the surface electrocardiogram, are shown to help predicting the outcome of catheter ablation of persistent atrial fibrillation. The proposed adaptive frequency tracking schemes are also applied to signals recorded in the field of sport sciences in order to illustrate their potential uses.

To summarize, the present thesis introduces several algorithms for adaptive frequency tracking. These algorithms are presented in full detail and they are then applied to practical situations. In particular, they are shown to improve the detection of coupling mechanisms in brain activity and to provide relevant organization measures for atrial fibrillation.

Keywords : instantaneous frequency, adaptive frequency tracking, adaptive filters, electroencephalography, cross-frequency couplings, electrocardiography, atrial fibrillation.

Résumé

L'amplitude et la fréquence sont les deux principaux attributs des signaux unidimensionnels, et ils sont donc tous deux utilisés pour l'analyse de données dans de nombreux domaines. Alors que l'amplitude peut être examinée directement, analyser la fréquence requiert des approches plus élaborées, excepté dans les cas les plus simples. Par conséquent, un grand nombre de techniques ont été proposées au fil des années pour extraire l'information fréquentielle. La méthode la plus connue est probablement l'estimation de la densité spectrale de puissance. Toutefois, cette approche est limitée aux signaux stationnaires puisque l'information temporelle est perdue. L'analyse temps-fréquence a été développée dans le but de résoudre le problème de l'estimation de fréquence pour des données non-stationnaires. Bien qu'elle permette d'estimer la puissance d'un signal dans un intervalle de temps donné et dans une gamme de fréquence données, ce type d'analyse a deux inconvénients qui diminuent son utilité dans certaines situations. Premièrement, à cause de l'interdépendance entre les résolutions temporelle et fréquentielle, améliorer la précision dans un domaine revient à la réduire dans l'autre. Deuxièmement, il est difficile d'utiliser ce type d'approche pour estimer la fréquence instantanée d'une composante oscillatoire spécifique. Une solution à ces deux limitations est apportée par les algorithmes adaptatifs de poursuite de fréquence. Ces algorithmes utilisent habituellement un filtre variant dans le temps (un filtre passe-bande ou un filtre coupe-bande) pour extraire une oscillation et un mécanisme adaptatif pour estimer sa fréquence instantanée.

La première partie de ce travail de thèse a pour objectif principal le développement d'un tel algorithme adaptatif de suivi de fréquence, le traqueur de fréquence. Les performances en termes de biais, de variance et de vitesse de convergence de ce dernier peuvent être comparées favorablement avec celles de méthodes existantes. La caractéristique la plus distinctive de cet algorithme adaptatif est qu'il tend à maximiser le comportement oscillatoire à sa sortie. En outre, grâce à un filtre spécifique, il n'introduit aucune distorsion dans la composante extraite. Des extensions pour cette méthode ont aussi été développées dans le but de prendre en compte certaines situations : la présence de plusieurs oscillations dans un signal, la question des composantes harmoniques, et la disponibilité de plus d'un signal avec la même oscillation. La première extension vise à suivre plusieurs composantes fréquentielles simultanément. L'idée de base est d'utiliser un traqueur pour estimer la fréquence de chaque oscillation. La deuxième extension utilise l'information contenue dans plusieurs signaux pour améliorer ses performances. Plus particulièrement, elle calcule séparément une estimée de la fréquence pour chaque signal disponible. Par la suite, ces estimées sont combinées avec des poids adaptatifs qui minimisent la variance. La troisième et dernière extension prend en compte une éventuelle structure harmonique dans un signal pour augmenter les performances de suivi. Cette extension est basée sur une idée similaire à la première et sur les mêmes poids que la deuxième. Une méthode itérative non-causale est aussi développée dans le but d'améliorer une trajectoire initiale de fréquence en utilisant l'information présente dans les échantillons futurs en plus de celle présente dans les échantillons passés. Comme le traqueur de fréquence, cette méthode cherche à maximiser le comportement oscillatoire à sa sortie. En revanche, elle est par définition uniquement applicable lorsque le traitement

des données est effectué en différé.

Dans la seconde partie de cette dissertation, les algorithmes de poursuite de fréquence développés dans la première partie sont appliqués à des données d'électroencéphalogramme et d'électrocardiogramme. Dans la première étude, le traqueur de fréquence est utilisée pour analyser les interactions ou couplages entre oscillations neuronales dans différentes gammes de fréquence. Plus précisément, l'évolution temporelle de ces couplages est examinée lors d'une expérience visuelle de potentiels évoqués avec des stimuli contenant des contours illusoires. Grâce à une approche adaptative qui garantit une information de phase correcte, les différences en termes de force de couplage entre stimuli avec et sans contours illusoires sont mises en évidence plus clairement qu'avec des méthodes traditionnelles basées sur des bancs de filtres prédéfinis. De plus, l'algorithme adaptatif permet de détecter des différences de fréquence dans certaines bandes. Dans une seconde étude, deux mesures d'organisation sont dérivées de l'extension pour composantes harmoniques. La première est basée sur la répartition de la puissance dans le domaine fréquentiel et la seconde sur la relation de phase entre harmoniques. Ces mesures, qui sont calculées à partir de l'électrocardiogramme de surface, aident à prédire le résultat d'une procédure d'ablation de la fibrillation auriculaire persistante. Les méthodes de suivi de fréquence proposées sont aussi appliquées à des signaux enregistrés dans le domaine du sport dans le but d'illustrer des utilisations potentielles.

Pour résumer, ce travail de thèse propose plusieurs techniques pour le suivi adaptatif de composantes fréquentielles. Ces technique sont d'abord présentées en détail, puis elles sont appliquées à des cas pratiques. En particulier, il est démontré qu'elles améliorent la détection des mécanismes de couplage dans l'activité du cerveau et qu'elles permettent de définir des mesures d'organisation pertinentes pour la fibrillation auriculaire.

Mot clés : fréquence instantanée, poursuite adaptative de fréquence, filtres adaptatifs, électroencéphalographie, couplages entre bandes de fréquence, électrocardiographie, fibrillation auriculaire.

Remerciements

Ce travail de thèse n’aurait pas été possible sans les contributions, directes ou indirectes, de nombreuses personnes. Je souhaite donc les remercier sincèrement :

- J’aimerais commencer par remercier les Profs. Reto A. Meuli et Micah M. Murray qui sont à l’origine de ce projet de recherche en collaboration entre le CHUV et l’EPFL. Grâce à leur implication et à leurs connaissances en neurosciences cette thèse a pu se dérouler dans les meilleures conditions.
- Merci au président du jury, le Prof. Pierre Vandergheynst, ainsi qu’aux rapporteurs, le Prof. Richard Reilly et les Drs. Paulo Gonçalves et Hervé Lissek, pour les discussions et commentaires lors de l’examen oral de cette thèse.
- Je tiens bien sûr à remercier les formidables “Jean-Marc Boys and Girls” que j’ai côtoyé au cours de ma thèse pour leur soutien, leur humour et leur bonne humeur : Laurent Uldry (alias mon Laurent, qui a pris le temps de relire ce manuscrit), Andrea Buttu, Cédric Duchêne, Yann ‘Lapin’ Prudat, Florian Jousset, Aline Cabasson, Alain Viso, Mathieu Lemay, Vincent Jacquemet et Meryem Jabloun. Merci aussi aux membres honoraires de l’ASPEGIC qui n’étaient jamais très loin : Elda, Guillaume, Benoît, Ashkan, Martin, Meri, Virginia (merci pour les cookies), et tous les autres. Vous avez toujours été là pour le travail, un conseil, un café ou toute autre activité.
- Plus généralement, j’aimerais remercier tous les membres des différents laboratoires de traitement du signal (ASPG, LTS et MMSPG) pour l’ambiance conviviale qui règne dans le couloir, ainsi que toutes les personnes que j’ai côtoyées durant ces quatre années de thèse. Un merci tout particulier à nos secrétaires Rosie De Pietro et Christine Gabriel pour leur efficacité et leurs sourires.
- Je souhaite aussi remercier les joueurs l’équipe de foot Galaxy qui est championne en titre du tournoi de l’EPFL. C’était toujours un plaisir de se changer les idées avec un petit match à midi.
- J’aimerais également exprimer toute ma gratitude à mes parents, Jocelyne et Alain, à mon petit frère, Hugo, et à toute ma famille pour leur soutien et leurs encouragements.
- Et finalement, un énorme merci à Jean-Marc Vesin, le légendaire patron de l’ASPG. Au fil des années, tu es devenu tellement plus qu’un directeur de thèse pour moi. Ton enthousiasme, tes connaissances encyclopédiques (pas uniquement en traitement du signal par ailleurs), ton humour incomparable ainsi que ton sens de la répartie inégalé ont rendu mes années de doctorat inoubliables. Je ne pourrai jamais te remercier assez pour tout ce que tu m’as apporté.

Contents

Abstract	v
Résumé	vii
Remerciements	ix
1 Introduction	1
1.1 Motivation and Problem Statement	1
1.2 Objectives	2
1.3 Organization	4
1.4 Original Contributions	5
I Frequency Estimation	7
2 Instantaneous Frequency Estimation	9
2.1 The Concept of Instantaneous Frequency	9
2.1.1 Analytic Signal	9
2.1.2 Instantaneous Frequency Definition	10
2.1.3 Interpretation of the Instantaneous Frequency	11
2.2 Methods for Instantaneous Frequency Estimation	11
2.2.1 Discrete Hilbert Transform	11
2.2.2 Energy Operators	14
2.2.3 Adaptive Line Enhancer	16
2.2.4 Oscillator-based Adaptive Band-pass Filter	18
2.2.5 Generalized Adaptive Notch Filter	19
2.2.6 Complex Adaptive Notch Filter	21
2.2.7 Robust Modified Newton Algorithm	22
2.3 Performance Comparison	24
2.4 Summary	27
3 Single Frequency Tracker	31
3.1 Introduction	31
3.2 Algorithm	31
3.3 Performance Analysis	33
3.3.1 Cost Function Analysis	35
3.3.2 Bias Analysis	38
3.3.3 Variance Analysis	43
3.4 Frequency Tracking Examples	47
3.5 Comparisons with Existing Tracking Algorithms	47

3.6	Discussion	52
4	Extensions for the Single Frequency Tracker	57
4.1	Introduction	57
4.2	Multiple Frequency Tracker	58
4.2.1	Multiple Frequency Tracking Example	59
4.3	Multivariate Frequency Tracker	60
4.3.1	Performance Analysis	64
4.3.2	Multivariate Multiple Frequency Tracker	66
4.3.3	Multivariate Frequency Tracking Examples	67
4.4	Harmonic Frequency Tracker	69
4.4.1	Performance Analysis	70
4.4.2	Harmonic Frequency Tracking Example	72
4.5	Discussion	72
5	Non-causal Instantaneous Frequency Enhancement	75
5.1	Introduction	75
5.2	Algorithm	76
5.3	Performance Analysis	79
5.4	Examples	83
5.5	Discussion	86
II	Frequency Tracking Applications	89
6	Adaptive Filters for the Identification of Couplings in Human EEG	91
6.1	Introduction	92
6.2	Materials and Methods	93
6.2.1	Experimental Setup	93
6.2.2	Data Acquisition and Pre-processing	94
6.2.3	Features	96
6.2.4	Synthetic Signals	98
6.3	Results	100
6.3.1	Synthetic Signals	100
6.3.2	Mean Frequency	101
6.3.3	Phase-amplitude Couplings	101
6.3.4	Phase-phase Couplings	104
6.3.5	Advantages of Adaptive Frequency Tracking	106
6.4	Discussion	111
7	Identification of the Outcome of Atrial Fibrillation Ablation	117
7.1	Introduction	117
7.2	Materials and Methods	119
7.2.1	Ablation Protocol	119
7.2.2	Patient Population and Electrophysiological Measurements	119
7.2.3	Signal Processing	119
7.2.4	Organization Measures	122
7.2.5	Statistical Analysis	123
7.3	Results	123
7.4	Discussion	126

8	Potential Applications of Adaptive Frequency Tracking	131
8.1	Heart Rate Estimation from Photoplethysmogram	131
8.2	Frequency Tracking and Near-infrared Spectroscopy	132
8.3	Indirect Estimation of the Respiration Frequency	135
8.4	Summary	137
9	Conclusion	139
9.1	Summary of Achievements	141
9.2	Perspectives	142
A	Notation and Definitions	145
A.1	Notation	145
A.2	Definitions	145
B	Performance Analysis of the Single Frequency Tracker	147
B.1	Input Signals	147
B.1.1	Noisy Cisoid	148
B.1.2	Noisy Sinusoid	148
B.1.3	Filtered Noisy Cisoid	149
B.1.4	Filtered Noisy Sinusoid	149
B.2	Cost Function Analysis	151
B.2.1	Noisy Cisoid	151
B.2.2	Noisy Sinusoid	151
B.2.3	Filtered Noisy Cisoid	153
B.2.4	Filtered Noisy Sinusoid	153
B.3	Bias Analysis	154
B.3.1	Noisy Cisoid	156
B.3.2	Noisy Sinusoid	156
B.3.3	Filtered Noisy Cisoid	156
B.3.4	Filtered Noisy Sinusoid	158
B.4	Variance Analysis	158
B.4.1	Noisy Cisoid	162
B.4.2	Noisy Sinusoid	162
B.4.3	Filtered Noisy Cisoid	163
B.4.4	Filtered Noisy Sinusoid	163
	Bibliography	169
	Index	183
	Curriculum Vitæ	185
	Publications	187

List of Acronyms

AF	atrial fibrillation
AM	amplitude modulation
DFT	discrete Fourier transform
DHT	discrete Hilbert transform
DTFT	discrete-time Fourier transform
ECG	electrocardiogram
EEG	electroencephalogram
FIR	finite impulse response
FM	frequency modulation
FT	Fourier transform
HFT	harmonic frequency tracker
HT	Hilbert transform
IIR	infinite impulse response
MFT	multiple frequency tracker
MSE	mean squared error
MVFT	multivariate frequency tracker
MVMFT	multivariate multiple frequency tracker
NFE	non-causal frequency estimator
NIRS	near-infrared spectroscopy
PPG	photoplethysmography
PSD	power spectral density
SFT	single frequency tracker
SNR	signal-to-noise ratio
STFT	short-time Fourier transform

List of Tables

2.1	Selected parameters for the performance comparison of the adaptive algorithms for instantaneous frequency estimation	26
2.2	Biases, variances and mean squared errors of the algorithms for instantaneous frequency estimation	27
2.3	Features of the algorithms for instantaneous frequency estimation	30
3.1	Default values and ranges for the parameters used in the performance analysis of the SFT	35
3.2	Selected parameters for the convergence speed comparisons of the adaptive frequency tracking algorithms	50
3.3	Statistics of the frequency tracking algorithms for the convergence speed comparisons	50
3.4	Selected parameters for the mean squared error comparisons of the adaptive frequency tracking algorithms for random walk frequency fluctuations	53
5.1	Mean squared errors	82
6.1	Results for the comparisons between real and surrogate datasets for the phase-amplitude and phase-phase couplings measured with the SFT	113
6.2	Results for the comparisons between IC and NC conditions for the phase-amplitude and phase-phase couplings measured with the SFT	114
6.3	Results for the comparisons between IC and NC conditions for the phase-amplitude and phase-phase couplings measured without the SFT	115
7.1	Patient characteristics	121

List of Figures

1.1	Two different non-stationary signals with the same power spectral density	3
1.2	Time-frequency representations of two different non-stationary signals with the same power spectral density	3
2.1	Example of problematic physical interpretation of instantaneous frequency	12
2.2	Distorsion between the impulse response of the Hilbert filter and its DFT approximation	14
2.3	Classical structure of adaptive frequency tracking filters	16
2.4	Frequency response of the adaptive line enhancer	17
2.5	Frequency response of the oscillator-based adaptive band-pass filter	19
2.6	Frequency responses of the notch and band-pass filters	22
2.7	Frequency estimation bias and variance of the DHT and DESA	25
2.8	Frequency responses of time-varying filters from ALE, OSC, GANF and CANF .	26
2.9	Convergence speeds of the algorithms for instantaneous frequency estimation . .	28
2.10	Convergence speeds of the GANF and CANF when applied to a real input signal and its analytic representation	28
3.1	Structure of the SFT	32
3.2	Band-pass filter response	32
3.3	Oscillator-based cost function for a noisy cisoid	37
3.4	Oscillator-based cost function for the analytic representation of a noisy sinusoid .	37
3.5	Oscillator-based cost function for a filtered noisy cisoid	39
3.6	Oscillator-based cost function for the filtered analytic representation of a noisy sinusoid	39
3.7	Bias of the SFT for a noisy cisoid	41
3.8	Bias of the SFT for the analytic representation of a noisy sinusoid	41
3.9	Bias of the SFT for a filtered noisy cisoid	41
3.10	Bias of the SFT for the filtered analytic representation of a noisy sinusoid	42
3.11	Variance of the SFT for a noisy cisoid	44
3.12	Variance of the SFT for the analytic representation of a noisy sinusoid	44
3.13	Variance of the SFT for a filtered noisy cisoid	45
3.14	Variance of the SFT for the filtered analytic representation of a noisy sinusoid . .	46
3.15	Instantaneous frequency estimated by the SFT	48
3.16	Input signal and oscillations extracted by the SFT	48
3.17	Instantaneous frequency estimated by the SFT for a linear chirp	49
3.18	Instantaneous frequency estimated by the SFT for FM oscillations	49
3.19	Convergence speeds of the frequency tracking algorithms for real- and complex-valued signals	51
3.20	Mean squared errors of the frequency tracking algorithms for random walk frequency fluctuations	54

4.1	Structure of the MFT	59
4.2	Global magnitude responses of the MFT	60
4.3	Instantaneous frequencies of three components estimated by the MFT	61
4.4	Input signal and three components extracted by the MFT	61
4.5	Structure of the MVFT	62
4.6	Validation of the proportionality assumption for the MVFT	64
4.7	Bias of the SFT and MVFT	65
4.8	Variance of the SFT and MVFT	65
4.9	Variance gain of the MVFT with respect to the SFT	65
4.10	Convergence speed of the SFT and MVFT	66
4.11	Instantaneous frequency and weights estimated by the MVFT	68
4.12	Instantaneous frequencies estimated by the MVMFT	68
4.13	Structure of the HFT	69
4.14	Bias of the SFT and HFT	71
4.15	Variance of the SFT and HFT	72
4.16	Convergence speed of the SFT and HFT	73
4.17	Instantaneous fundamental frequency of two FM harmonic components estimated by the HFT	74
4.18	Input signal and the two FM harmonic components extracted by the HFT	74
5.1	Instantaneous frequencies in the two considered cases for analyzing the NFE performance	80
5.2	Mean squared errors for the constant frequency case	81
5.3	Mean squared errors for the linear chirps case	81
5.4	Convergence speed of the SFT and NFE	83
5.5	Frequency estimation for a frequency modulated cisoid embedded in white noise	84
5.6	Frequency estimation for a complex linear chirp embedded in white noise	85
5.7	Estimation of the instantaneous frequency in an EEG signal	87
5.8	Estimation of the instantaneous frequency in an R-R interval signal	87
6.1	Processing steps used for analyzing the EEG data from the IC experiment	94
6.2	Experimental conditions	94
6.3	Pre-processing steps	96
6.4	Original and surrogate EEG signals	97
6.5	Examples of the synthetic signals used for assessing the usefulness of the SFT	100
6.6	Mean PLV values obtained with and without adaptive frequency tracking	102
6.7	Mean PLV values obtained with and without adaptive frequency tracking for a synthetic EEG example	102
6.8	Comparisons between real and surrogate datasets for the mean estimated frequency	103
6.9	Comparisons between IC and NC conditions for the mean estimated frequency	103
6.10	Comparisons between real and surrogate datasets for the phase-amplitude couplings	105
6.11	Comparisons between IC and NC conditions for the phase-amplitude couplings	105
6.12	Comparisons between real and surrogate datasets for the phase-phase couplings	107
6.13	Comparisons between IC and NC conditions for the phase-phase couplings	107
6.14	Comparisons between real and surrogate datasets for phase-amplitude and phase-phase couplings measured with the SFT	108
6.15	Comparisons between IC and NC conditions for phase-amplitude and phase-phase couplings measured with the SFT	109
6.16	Comparisons between IC and NC conditions for phase-amplitude and phase-phase couplings measured without the SFT	110

7.1	Stepwise radiofrequency catheter ablation protocol	120
7.2	Harmonic frequency tracking example	124
7.3	Example of real signal of atrial activity	125
7.4	Pearson’s correlation coefficients between chest leads and intracardiac AFCLs at baseline	125
7.5	Boxplots of the four organization measures at baseline with respect to the three groups	127
7.6	Means and standard deviations of the four organization measures at baseline with respect to the three groups	128
8.1	PPG signal recorded from the knee	133
8.2	Heart rate estimated from a PPG signal	133
8.3	Excerpt of an O2Hb signal during sprint interval training	134
8.4	Estimation of the pedaling frequency from NIRS data	136
8.5	Respiration component extracted from R-R intervals	138
8.6	Respiration frequency during an incremental exercise test	138
B.1	Oscillator-based cost function for a noisy cisoid	152
B.2	Oscillator-based cost function for the analytic representation of a noisy sinusoid	152
B.3	Oscillator-based cost function for a filtered noisy cisoid	155
B.4	Oscillator-based cost function for the filtered analytic representation of a noisy sinusoid	155
B.5	Bias of the SFT for a noisy cisoid	157
B.6	Bias of the SFT for the analytic representation of a noisy sinusoid	157
B.7	Bias of the SFT for a filtered noisy cisoid	159
B.8	Bias of the SFT for the filtered analytic representation of a noisy sinusoid	159
B.9	Variance of the SFT for a noisy cisoid	164
B.10	Variance of the SFT for the analytic representation of a noisy sinusoid	164
B.11	Variance of the SFT for a filtered noisy cisoid	165
B.12	Variance of the SFT for the filtered analytic representation of a noisy sinusoid	167

1

Introduction

1.1 Motivation and Problem Statement

In many applications, most of the information carried in signals is contained in their amplitude and their frequency content. Typically, these characteristics are examined from two complementary viewpoints: the time and frequency domains. These two domains are alternative perspectives on the power of a signal. The former is particularly suited for investigating its temporal evolution, whereas the latter is mostly relevant when inspecting its repartition among the frequencies. As both temporal and frequency power distributions are key elements in numerous fields, a large number of techniques have been developed over the years to estimate these features. The amplitude information can be straightforwardly observed from the considered signal, and the temporal distribution of power can be easily investigated with simple tools, such as local estimates of the variance. Naturally, more complex methods have been proposed with desired properties, such as noise resilience or estimation robustness. On the other hand, except in very simple cases, the information about frequency is not directly available. Usually, the signal of interest is first transformed in the frequency domain before carrying out the observation. The most widely used technique to analyze the frequency content of a signal is power spectral density (PSD) estimation [1]. Many different approaches have been developed for this purpose and they can be classified in two categories: non-parametric and parametric methods. The non-parametric ones do not make any assumption about the considered signal and are applicable to any data without restrictions. The best known techniques in this category are the periodogram, the averaged periodogram and the modified periodogram [2, 3]. On the other hand, the parametric approaches assume that the signal under study has a certain stochastic structure which can be used to improve the accuracy and robustness of the estimation with respect to non-parametric methods. It is clear that better estimation performance is only achieved when the assumptions are verified to a certain extent. Otherwise, parametric approaches may perform worse than the non-parametric ones. Well-known parametric methods assume, for instance, that the considered signal can be closely approximated with autoregressive or autoregressive moving average models [4], or has a line spectrum [5–7].

A very important requirement of all approaches for PSD estimation is that the signal under study has to be stationary. Even non-parametric techniques are based on this assumption. Indeed, as the temporal information is lost in a PSD estimate, the statistical properties of the considered signal should not change over time. Strict stationarity is not necessarily required in most cases. Indeed, a weaker concept, known as wide-sense stationarity [8], is typically used instead. The first and second moments of a wide-sense stationary stochastic process do not vary with respect

to time. As long as this assumption holds, PSD estimation techniques provide accurate and reliable assessment of the frequency content of a signal. It is worth mentioning that, for some methods, the stationary assumption needs only to be valid up to a certain extent. However, for non-stationary data, the PSD approaches are not sufficient as temporal information is lost. In fact, different non-stationary signals can have the same PSD. This issue is illustrated in Figure 1.1 for two sinusoids whose instantaneous frequencies vary linearly over time, respectively from 0.2π to 0.4π and 0.4π to 0.2π . Although they are different, their PSDs are identical. This is confirmed by their PSD estimates, computed with Welch's method [2], that overlap perfectly. Furthermore, these estimates give the false idea that there is power in the range from 0.2π to 0.4π for the whole duration of the signals, while there is only a single periodic component with time-varying frequency. In order to cope with non-stationary data, many different time-frequency analyses were developed [9–12]. With these techniques, the power at given time and frequency can be evaluated. The two main approaches to obtain time-frequency representations are local Fourier transforms and wavelet analyses. The most famous method is probably the short-time Fourier transform (STFT), wherein the discrete Fourier transform (DFT) of the signal of interest is computed over sliding windows. Figure 1.2 shows the STFTs of the two sinusoids used to emphasize the issue with PSD estimation of non-stationary data. Owing to the time-frequency representations, the temporal evolution of the frequency content of both signals is clearly visible.

Despite solving the temporal information loss problem of PSD estimation, time-frequency approaches are not the ultimate panacea. They have several drawbacks that can become too penalizing in certain situations. First, with the exception of the simplest cases such as the one presented in Figures 1.1 and 1.2, it is difficult to retrieve the instantaneous frequency of an oscillatory component from a time-frequency representation. Indeed, such techniques do not provide this kind of information directly and therefore further processing is required. In the presence of random noise and interfering oscillations, estimating the instantaneous frequency can be a very difficult task. Furthermore, if the component corresponding to the estimated frequency is needed, an additional processing step is required, as no extraction is performed by default. Another limitation is that time-frequency representations are computed block by block. This has an impact on the temporal accuracy of the estimation, especially during transition periods. A closely related issue is the question of time and frequency resolutions which is inherent to this kind of analysis. Indeed, increasing the time resolution decreases the frequency resolution and vice versa. Therefore, there is a tradeoff between the two resolutions. In addition, although computational power is not really an issue anymore due to advances in computer technology, the computational load of time-frequency analyses may become prohibitive in certain applications, such as in embedded devices or during very large-scale investigations. Collectively, the limitations of PSD estimation and time-frequency techniques call for efficient schemes to estimate the instantaneous frequency and extract the underlying oscillatory component.

1.2 Objectives

The present dissertation is aimed at the development of an algorithm for tracking a single oscillatory component and estimating its instantaneous frequency with several desirable properties, namely accuracy, reliability, robustness to disturbances, applicability to various fields, and low complexity. Related objectives are to extend this adaptive scheme to three specific cases of practical interest: multiple frequency tracking, multivariate frequency tracking and harmonic frequency tracking. The first extension is capable to handle signals that contain more than one oscillatory components. The two other extensions improve the overall tracking performance by taking advantage of additional information provided by several signals or by taking into account an harmonic structure. Another goal is to develop a non-causal method for enhancing frequency trajectories. As data analysis is often carried out offline, the frequency estimation performance

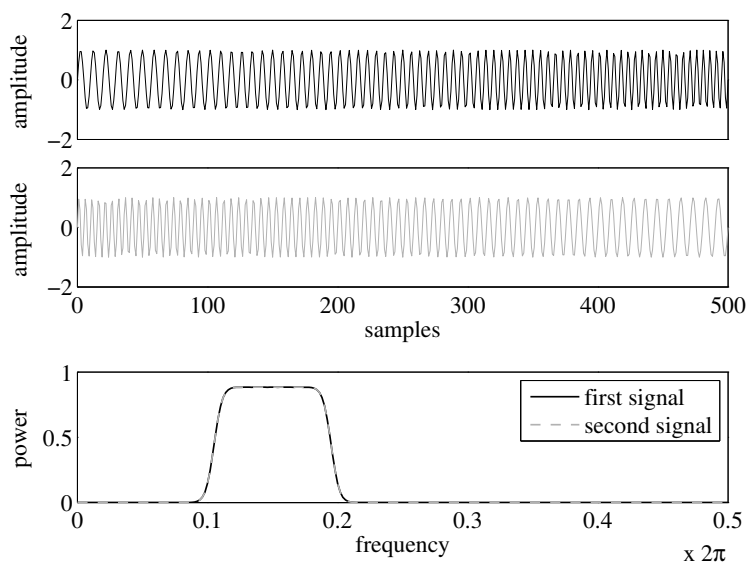


Figure 1.1: Two different non-stationary signals with the same power spectral density (PSD). From top to bottom are plotted the first and second signals, and their PSD estimates. Both estimates overlap perfectly.

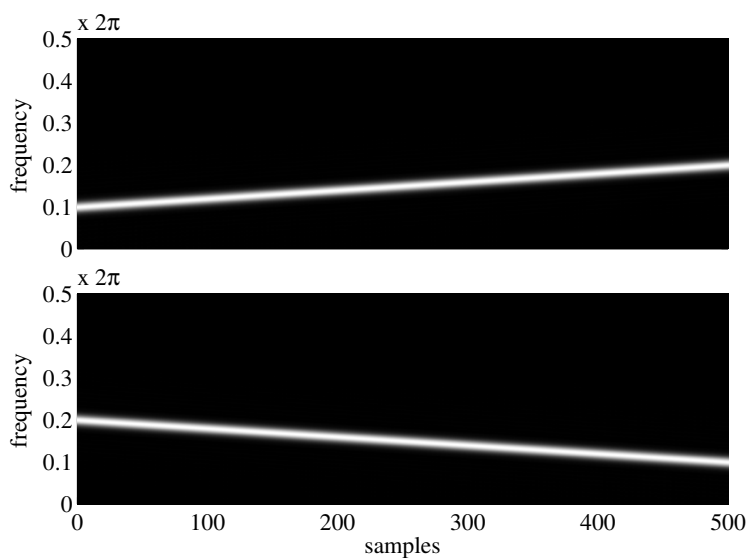


Figure 1.2: Time-frequency representations of two different non-stationary signals (plotted in Figure 1.1) with the same power spectral density. These representations were computed using the short-time Fourier transform.

can be improved by using future information in addition to past information. The last two objectives are to show the relevance of the algorithms developed during this thesis for processing electroencephalographic (EEG) and electrocardiographic (ECG) data. Specifically, adaptive frequency tracking is applied to highlight cross-frequency coupling mechanisms in EEG signals recorded during a visual evoked potential experiment and to extract organization measures from ECG data in order to predict the outcome of stepwise catheter ablation of atrial fibrillation. To summarize, the main objectives of this dissertation are:

- to develop an efficient scheme for adaptive frequency tracking,
- to extend this scheme in order to track several oscillations simultaneously,
- to improve the overall tracking performance by taking advantage of multivariate data, or by taking into account the presence of an harmonic structure,
- to enhance frequency trajectories with non-causal processing,
- to apply the developed algorithms to EEG data in order to highlight cross-frequency coupling mechanisms taking place in the brain during visual perception,
- to extract organization measures with adaptive frequency tracking from ECG data for predicting the outcome of atrial fibrillation ablation.

This work has been realized in the framework of a collaboration between the Applied Signal Processing Group¹ of the Swiss Federal Institute of Technology and the Departments of Clinical Neurosciences and Radiology of the Centre Hospitalier Universitaire Vaudois.² The ultimate goal of this thesis is to provide efficient schemes for adaptive frequency tracking that are relevant for, but restricted to, biomedical data, such as EEG and ECG.

1.3 Organization

This dissertation is divided into two parts. The first part introduces techniques for estimating the instantaneous frequency, with a particular emphasis on adaptive frequency tracking. In the second part, applications of these methods to real data, such as EEG and ECG signals, are presented.

I Frequency estimation. After introducing the concept of instantaneous frequency, a selection of techniques for its estimation are presented in Chapter 2. In Chapter 3, the single frequency tracker, an adaptive frequency tracking scheme developed during this thesis, is detailed and its performance is analyzed thoroughly. This technique is extended to take into account multiple oscillatory components, multiple signals and harmonic frequencies in Chapter 4. To conclude this part, an iterative algorithm for non-causal enhancement of frequency trajectories is developed in Chapter 5.

II Frequency tracking applications. Chapter 6 presents a study in which adaptive frequency tracking was applied to EEG signals with the purpose of identifying interactions or couplings between neuronal oscillatory components in different frequency bands. In Chapter 7, the harmonic extension of the single frequency tracker is shown to help predicting the outcome of catheter ablation of atrial fibrillation. At this end of this part, three examples of applications for the presented frequency tracking schemes are given in Chapter 8.

Finally, Chapter 9 concludes this dissertation by providing a summary of the results achieved during this thesis as well as possible future research directions. It is also worth to mention that the notation used throughout this dissertation is given in Appendix A alongside a few relevant definitions.

1. <http://aspg.epfl.ch/>

2. <http://www.chuv.ch/>

1.4 Original Contributions

The main contributions³ of this work are:

- Development of efficient adaptive schemes for instantaneous frequency estimation.
 - The single frequency tracker, an algorithm for tracking a single oscillatory component.
 - Extensions of this algorithm for three specific cases:
 - Multiple frequency tracking.
 - Multivariate frequency tracking.
 - Harmonic frequency tracking.
 - A non-causal iterative technique for enhancing frequency trajectories.
- Applications to biomedical data.
 - The use of the single frequency tracker to highlight cross-frequency couplings between brain oscillations during a visual evoked potential experiment.
 - The use of harmonic frequency tracking to develop organization measures for predicting the success of catheter ablation in patients with persistent atrial fibrillation.

3. See also the list of publications at the end of the text (p. 187).

Part I

Frequency Estimation

Instantaneous Frequency Estimation

2

Frequency information is a key aspect in a wide range of fields. The instantaneous frequency, which is typically defined as the frequency at each time, can shed light on a large number of non-stationary phenomena. It is therefore useful in many applications such as communications [13, 14], speech processing [15–17] and biomedical signal processing [18, 19]. However, accurately estimating the instantaneous frequency is not straightforward except in very simple situations, and thus many different methods have been proposed over the years for this purpose.

The concept of instantaneous frequency is reviewed in Section 2.1. Then, Section 2.2 presents a selection of methods for instantaneous frequency estimation. An evaluation of their performance is provided in Section 2.3. Finally, the strengths and weaknesses of the presented methods are summarized in Section 2.4.

2.1 The Concept of Instantaneous Frequency

Frequency is defined as the number of oscillations observed during one unit of time. By analogy with this definition, the instantaneous frequency is the frequency at a given time. However, it is not applicable to all types of signals. In addition, its interpretation is controversial. A general introduction to the concept of instantaneous frequency and its interpretation is provided here. For a more thorough presentation see [20].

Carson and Fry [21] first introduced the concept of instantaneous frequency with complex frequency-modulated (FM) signals. They defined it as the rate of change of phase angle at a given time. It is therefore a generalization of the definition of constant frequency. Later, van der Pol [22] also investigated this concept. He tackled this issue by considering amplitude- and frequency-modulated sinusoids with the following form:

$$x(t) = a(t) \cos(\phi(t)), \quad (2.1)$$

where $a(t)$ and $\phi(t)$ are the time-varying amplitude and phase respectively. Based on this harmonic expression, he defined the instantaneous frequency as the derivative of the phase with respect to time:

$$\omega_i(t) = \frac{d\phi(t)}{dt}. \quad (2.2)$$

2.1.1 Analytic Signal

Further study of instantaneous frequency required a method for extracting the phase from all types of signals and not only sinusoids. Gabor [23] made it possible by proposing a method for

uniquely generating a complex signal, known as the analytic signal, from a real one.

The procedure for computing this analytic representation is to first take the Fourier transform (FT) of the real signal. Then, the amplitudes of negative frequencies are set to zero while the amplitudes of positive frequencies are multiplied by two. Finally, the modified spectrum is transformed back to the time domain in order to obtain the analytic signal. This whole procedure in the frequency domain is equivalent to the following time domain approach,

$$x_a(t) = x(t) + j\mathcal{H}\{x(t)\} = x(t) + jx_h(t) = a(t)e^{j\phi(t)}. \quad (2.3)$$

In this expression $x_a(t)$ is the analytic signal, $x(t)$ is the real signal, $\mathcal{H}\{\cdot\}$ denotes the Hilbert transform (HT). Thus, $x_h(t)$ is the Hilbert transformed original signal. Moreover, $a(t)$ and $\phi(t)$ are the instantaneous amplitude and phase. The HT is defined as

$$x_h(t) = \mathcal{H}\{x(t)\} = \frac{1}{\pi} \text{p.v.} \int_{-\infty}^{\infty} \frac{x(\tau)}{t - \tau} d\tau, \quad (2.4)$$

where p.v. denotes the Cauchy principal value of the integral. Due to the integration interval, knowledge of $x(t)$ over the whole time domain is required for computing the HT. It is thus non-local. Nonetheless, the main contribution of the integral (2.4) is made in the neighborhood of time t .

The HT can also be seen as the convolution of $x(t)$ with the function $\frac{1}{\pi t}$. Therefore, in the Fourier domain, the HT is expressed as

$$X_h(f) = \mathcal{F} \left\{ \frac{1}{\pi t} \right\} \cdot X(f). \quad (2.5)$$

Using the following FT pair [24],

$$\frac{1}{t} \xleftrightarrow{\mathcal{F}} -j\pi \operatorname{sgn}(f), \quad (2.6)$$

$X_h(f)$ can be expressed as

$$X_h(f) = \begin{cases} jX(f) & \text{for } f < 0, \\ 0 & \text{for } f = 0, \\ -jX(f) & \text{for } f > 0. \end{cases} \quad (2.7)$$

Therefore, the imaginary part of the analytic signal can be seen as the original signal delayed by a phase of $-\frac{\pi}{2}$. Based on this result, the analytic signal of an harmonic oscillation $x(t) = A_0 \cos(\omega_0 t)$ is $x_a(t) = A_0 \cos(\omega_0 t) + jA_0 \sin(\omega_0 t)$. So the HT performs a phase shift of $-\frac{\pi}{2}$ (respectively $\frac{\pi}{2}$) for every spectral components with positive (respectively negative) frequencies.

2.1.2 Instantaneous Frequency Definition

Based on the work of Carson and Fry [21], van der Pol [22] and Gabor [23], Ville [25] defined the instantaneous frequency as the derivative of the instantaneous phase $\phi(t)$ of the analytic signal $x_a(t)$ with respect to time:

$$\omega_i(t) = \frac{d}{dt} \arg\{x_a(t)\}. \quad (2.8)$$

This is the most frequently used definition.

2.1.3 Interpretation of the Instantaneous Frequency

Despite the extended usage of the concept of instantaneous frequency, its definition as well as its interpretation are still subject to controversy. For example, the definition (2.8) is not a unique function of time for Shekel [26]. Indeed, any complex amplitude-modulated (AM) wave can be written either as $a(t)e^{j\omega_0 t}$ or as $a_0 e^{j\phi(t)}$. This means that there exist many different complex representations of a real signal. Moreover, although a unique complex representation of any real signal can be obtained with the HT (2.4), there is no certainty that the instantaneous frequency (2.8) corresponds to a physical reality. In fact, the instantaneous amplitude $a(t)$ and instantaneous phase $\phi(t)$ have a clear physical meaning for narrow-band signals only [27]. Furthermore, Mandel [28] argued that there is no one-to-one relationship between the instantaneous frequency and frequencies obtained with Fourier decomposition, and therefore challenged any physical interpretation. He used the following signal to illustrate his viewpoint:

$$x(t) = a_1 e^{j(\omega_0 - \frac{1}{2}\omega_\Delta)t} + a_2 e^{j(\omega_0 + \frac{1}{2}\omega_\Delta)t} = a(t)e^{j\phi(t)} \quad (2.9)$$

with

$$a(t) = \sqrt{a_1^2 + a_2^2 + 2a_1 a_2 \cos(\omega_\Delta t)},$$

$$\phi(t) = \omega_0 t + \arctan\left(\frac{(-a_1 + a_2) \sin(\frac{1}{2}\omega_\Delta t)}{(a_1 + a_2) \cos(\frac{1}{2}\omega_\Delta t)}\right).$$

The instantaneous frequency of $x(t)$ is given by

$$\omega_i(t) = \omega_0 + \frac{\omega_\Delta}{2} \frac{-a_1^2 + a_2^2}{a_1^2 + a_2^2 + 2a_1 a_2 \cos(\omega_\Delta t)}. \quad (2.10)$$

An interpretation problem arises in (2.10) if a_1 is not equal to a_2 . In this case there is no symmetry in the fluctuations of instantaneous frequency around ω_0 . Moreover, $\omega_i(t)$ is always below ω_0 when $a_1 > a_2$ and always above when $a_1 < a_2$. This problem is illustrated in Figure 2.1. This missing relationship between the instantaneous frequency and the average frequency computed from the instantaneous spectrum as well as the resulting lack of physical interpretation have been highlighted by many authors [29]. In fact, the instantaneous frequency corresponds to the average frequency at each time only when there is symmetry in the instantaneous spectrum [30]. It is also important to mention that some authors consider signals with such problematic physical interpretation as multi-component ones and thus outside of the definition of instantaneous frequency [20].

2.2 Methods for Instantaneous Frequency Estimation

A large number of methods for estimating the instantaneous frequency have been proposed over the years. Indeed, various starting points have been considered in order to tackle this estimation problem. In this section, a small selection of methods are presented. First, two instantaneous techniques are described: the classical method based on the discrete Hilbert transform and an algorithm relying on an energy operator. Then, four adaptive schemes using time-varying band-pass (or notch) filters are introduced. Finally, a technique based on linear prediction is presented.

2.2.1 Discrete Hilbert Transform

The instantaneous frequency of a signal can be obtained through its analytic representation as presented in Section 2.1 in the continuous-time case. This approach is also applicable to

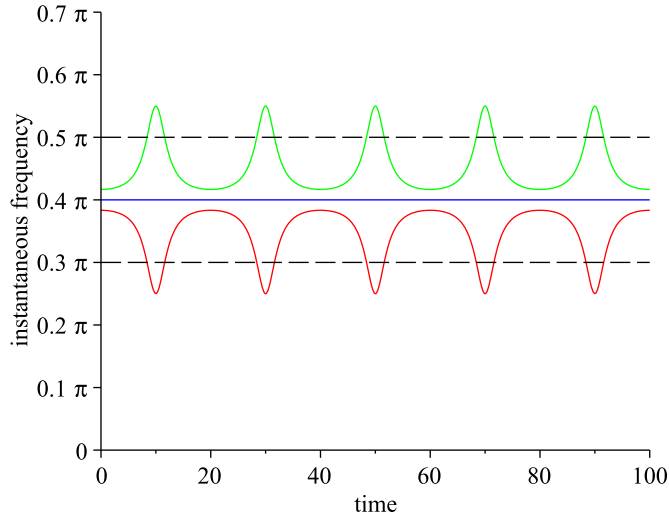


Figure 2.1: Example of problematic physical interpretation of the instantaneous frequency for signal (2.9) with $\omega_0 = 0.4\pi$ and $\omega_\Delta = 0.1\pi$ for the three cases $a_1 = a_2$ (blue), $a_1 > a_2$ (red) and $a_1 < a_2$ (green). The two Fourier components of the signal are denoted by the dashed black lines.

discrete-time signals. In this case, the analytic signal is defined as

$$x_a[n] = x[n] + jx_h[n] = a[n]e^{j\phi[n]}, \quad (2.11)$$

where $x_h[n]$ is the discrete Hilbert transform (DHT) of $x[n]$. As for the continuous case, the DHT is computed by convolving a signal with the Hilbert filter [31]: $x_h[h] = h[n] * x[n]$. However, the discrete-time filter is defined slightly different from its continuous-time counterpart (2.4). This infinite impulse response (IIR) filter is given by

$$h[n] = \begin{cases} 0 & \text{for } n \text{ even,} \\ \frac{2}{\pi n} & \text{for } n \text{ odd.} \end{cases} \quad (2.12)$$

The DHT is square summable for square summable signals [32]. In practical situations, the Hilbert filter (2.12) is approximated with a finite impulse response (FIR). The DHT can also be computed in the frequency domain [33]. In this case, the convolution becomes a multiplication:

$$X_h(e^{j\omega}) = \text{DTFT}\{x_h[n]\} = \text{DTFT}\{h[n] * x[n]\} = H(e^{j\omega}) \cdot X(e^{j\omega}), \quad (2.13)$$

where the discrete-time Fourier transform (DTFT) of the Hilbert filter (2.12) is defined as

$$H(e^{j\omega}) = \begin{cases} j & \text{for } -\pi < \omega < 0, \\ 0 & \text{for } \omega = 0, \\ -j & \text{for } 0 < \omega < \pi. \end{cases} \quad (2.14)$$

Therefore, the DHT is obtained simply by shifting the phases of positive frequency components by $-\frac{\pi}{2}$ and the ones of negative frequency components by $\frac{\pi}{2}$.

Since the DTFT is not applicable in practice, the discrete Fourier transform (DFT) is used

instead. In this case, the frequency response of the Hilbert filter of length N is given by

$$H[k] = \begin{cases} 0 & \text{for } k = 0, \\ -j & \text{for } 1 \leq k \leq \frac{N}{2} - 1, \\ 0 & \text{for } k = \frac{N}{2}, \\ j & \text{for } \frac{N}{2} + 1 \leq k \leq N - 1, \end{cases} \quad \text{when } N \text{ is even,} \quad (2.15a)$$

and

$$H[k] = \begin{cases} 0 & \text{for } k = 0, \\ -j & \text{for } 1 \leq k \leq \frac{N-1}{2}, \\ j & \text{for } \frac{N+1}{2} \leq k \leq N - 1, \end{cases} \quad \text{when } N \text{ is odd.} \quad (2.15b)$$

The problem is that the sampling in the frequency domain introduced by the transition from the DTFT to the DFT leads to some distortion. Indeed, the DFT only gives an approximation of the true impulse response of the Hilbert filter (2.12). The agreement between the true filter and its DFT approximation depends on the signal length as illustrated in Figure 2.2. Nevertheless, the amount of distortion is negligible when a sufficient number of samples are available, and difficulties should occur only for signals with very short lengths. It is also important to note that typical implementations for computing the analytic signal from the DFT skip the DHT. Instead, $X_a[k]$ is computed directly from the DFT of a real signal of length N [34]:

$$X_a[k] = \begin{cases} X[0] & \text{for } k = 0, \\ 2X[k] & \text{for } 1 \leq k \leq \frac{N}{2} - 1, \\ X[N/2] & \text{for } k = \frac{N}{2}, \\ 0 & \text{for } \frac{N}{2} + 1 \leq k \leq N - 1, \end{cases} \quad \text{when } N \text{ is even,} \quad (2.16a)$$

and

$$X_a[k] = \begin{cases} X[0] & \text{for } k = 0, \\ 2X[k] & \text{for } 1 \leq k \leq \frac{N-1}{2}, \\ 0 & \text{for } \frac{N+1}{2} \leq k \leq N - 1, \end{cases} \quad \text{when } N \text{ is odd.} \quad (2.16b)$$

Once the analytic signal has been computed, the instantaneous phase is easily extracted by taking its argument,

$$\phi[n] = \arg\{x_a[n]\}. \quad (2.17)$$

Then, the instantaneous frequency is obtained by differentiating the phase. The finite differences methods are a very simple yet efficient approach for performing the discrete-time differentiation. In particular, forward, backward and central finite differences can be used for computing the instantaneous frequency:

$$\omega_f[n] = \phi[n+1] - \phi[n], \quad (2.18a)$$

$$\omega_b[n] = \phi[n] - \phi[n-1], \quad (2.18b)$$

$$\omega_c[n] = \frac{\phi[n+1] - \phi[n-1]}{2}. \quad (2.18c)$$

More complex methods with desirable properties (e.g. noise resilience) are also available to perform the differentiation operation [35]. It is also worth to mention that, in addition to the instantaneous phase, the HT also provides the envelope of the original signal as its instantaneous amplitude.

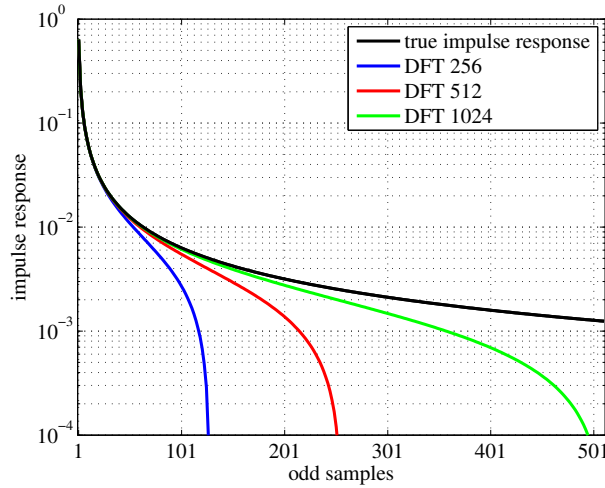


Figure 2.2: Distorsion between the impulse response of the Hilbert filter (2.12) and its DFT approximation for various lengths. Only the positive odd samples are shown because the even samples are equal to zero and the impulse response is antisymmetric.

2.2.2 Energy Operators

In his work on nonlinear speech modeling, Teager [36–38] developed a simple yet effective operator Ψ for tracking the energy present in a speech signal. It is based on the observation that speech signals should be analyzed from the point of view of the energy required to generate them. This energy operator is defined as

$$\Psi_c\{x(t)\} = \dot{x}^2(t) - x(t)\ddot{x}(t) \quad (2.19)$$

with $\dot{x}(t) = dx(t)/dt$ and $\ddot{x}(t) = d^2x(t)/dt^2$ for continuous-time data, and as

$$\Psi_d\{x[n]\} = x^2[n] - x[n+1]x[n-1] \quad (2.20)$$

for discrete-time data. Kaiser [39], who introduced both continuous- and discrete-time forms systematically, showed that this operator can track the energy of a linear oscillator.

Furthermore, under certain conditions, the energy operator Ψ can be used to extract the instantaneous amplitude and frequency from amplitude- and frequency-modulated (AM-FM) signals [40]. In continuous time, an AM-FM signal is defined as

$$x(t) = A(t) \cos(\phi(t)) = A(t) \cos\left(\omega_c t + \omega_m \int_0^t q(\tau) d\tau + \theta\right), \quad (2.21)$$

where $A(t)$ is the amplitude, ω_c is the carrier frequency, ω_m is the maximum deviation from ω_c , $q(t)$ is the modulating (or information) signal, and $\theta = \phi(0)$ is an arbitrary phase offset. These variables are subject to the following constraints:

$$-1 \leq q(t) \leq 1, \quad \forall t \quad \text{and} \quad 0 < \omega_m < \omega_c. \quad (2.22)$$

The instantaneous frequency is computed as follows,

$$\omega_i(t) = \frac{d\phi(t)}{dt} = \omega_c + \omega_m q(t). \quad (2.23)$$

Applying the continuous form of the operator to $x(t)$ (2.21) yields the squared product of the instantaneous amplitude and frequency:

$$\Psi_c\{x(t)\} \approx (A(t)\omega_i(t))^2. \quad (2.24)$$

Similarly, an AM-FM signal in discrete time is given by

$$x[n] = A[n] \cos(\phi[n]) = A[n] \cos\left(\omega_c n + \omega_m \int_0^n q[k] dk + \theta\right), \quad (2.25)$$

and its instantaneous frequency is

$$\omega_i[n] = \frac{d\phi[n]}{dn} = \omega_c + \omega_m q[n]. \quad (2.26)$$

All parameters are defined in the same way as for the continuous case. Also, in the two previous equations, both the integration $\int dk$ and differentiation d/dn treat the discrete-time indices k and n symbolically as continuous variables. Applying the discrete form of the operator as before yields a slightly different result for a discrete-time AM-FM signal:

$$\Psi_d\{x[n]\} \approx A^2[n] \sin^2(\omega_i[n]). \quad (2.27)$$

Therefore, based on the results in continuous and discrete time, the energy operator Ψ can only extract the amplitude and frequency as a product, but it cannot separate them. A solution to this problem was proposed by Maragos et al. [41]. In continuous time, the idea is to use, in addition to the signal, its time derivative. Thus, the instantaneous frequency and amplitude are obtained as

$$\omega_i(t) \approx \sqrt{\frac{\Psi_c\{\dot{x}(t)\}}{\Psi_c\{x(t)\}}} \quad (2.28)$$

and

$$|A(t)| \approx \frac{\Psi_c\{x(t)\}}{\sqrt{\Psi_c\{\dot{x}(t)\}}}. \quad (2.29)$$

This method is called the continuous energy separation algorithm. In discrete time, this approach translates into using the difference signal, which is defined as

$$y[n] = x[n] - x[n-1]. \quad (2.30)$$

The instantaneous frequency and amplitude are then extracted as follows,

$$\omega_i[n] \approx \arccos\left(1 - \frac{\Psi_d\{y[n]\} + \Psi_d\{y[n+1]\}}{4\Psi_d\{x[n]\}}\right) \quad (2.31)$$

and

$$|A(n)| \approx \sqrt{\frac{\Psi_d\{x[n]\}}{1 - \left(1 - \frac{\Psi_d\{y[n]\} + \Psi_d\{y[n+1]\}}{4\Psi_d\{x[n]\}}\right)^2}}. \quad (2.32)$$

This technique is known as the discrete energy separation algorithm.

2.2.3 Adaptive Line Enhancer

The majority of adaptive frequency tracking filters, also known as adaptive line enhancers, have a similar structure composed of two main parts: a time-varying band-pass filter and an adaptive mechanism. This structure is illustrated in Figure 2.3. The purpose of the filter, which is replaced by a time-varying notch filter in some algorithms, is to extract or enhance (hence the term ‘adaptive line enhancer’) the periodic component in the input signal. On the other hand, the adaptive mechanism estimates the instantaneous frequency on the basis of the output signal and, in turn, updates the filter through an adaptive parameter. This parameter typically reflects the current frequency estimate through some mathematical operations. As numerous applications require reliable estimates of the instantaneous frequency, many different adaptive frequency tracking filters (or adaptive line enhancers) have been proposed over the years [42–48]. In these algorithms, the periodic component is typically extracted with either FIR [43, 44] or IIR [45–48] filters. The well-known FIR filters are inherently stable and easy to adapt. However, they can be very long as many coefficients are required to reach a sufficiently narrow bandwidth. By contrast, narrow IIR filters can have an order as small as two (or even one for complex-valued signals). Also, some of these filters can be adapted efficiently through only one parameter, as long as the poles are constrained inside the unit circle in order to ensure stability. The adaptive mechanism that updates the current frequency estimate is usually designed to minimize the mean squared error between input and output signals.

To illustrate this class of algorithms for instantaneous frequency estimation, the adaptive line enhancer proposed by Hush et al. [49] is presented in more detail. The input signal is defined as

$$x[n] = s[n] + v[n] = A_0 \sin(\omega_0 n + \phi_0) + v[n], \quad (2.33)$$

where A_0 , ω_0 and ϕ_0 are respectively the amplitude, frequency and phase of the sinusoid, and $v[n]$ is an additive white noise. In practice, the amplitude and frequency can change over time. Nonetheless, as the presented algorithm is adaptive, it is capable to adjust to these changes. The periodic component is extracted with a time-varying band-pass filter whose transfer function is given by

$$H(z; n) = \frac{\frac{1-r^2}{1+r^2} \alpha[n] z^{-1} - (1-r^2) z^{-2}}{1 - \alpha[n] z^{-1} + r^2 z^{-1}}, \quad (2.34)$$

where r ($0 \ll r < 1$) determines the radius of the poles and thus the bandwidth, and $\alpha[n]$ ($-2r < \alpha[n] < 2r$) is the adaptive parameter at time n . It is linked to the current frequency estimate $\omega[n]$, which is also the central frequency of the filter (2.34), through the following relation:

$$\omega[n] = \arccos\left(\frac{\alpha[n]}{1+r^2}\right). \quad (2.35)$$

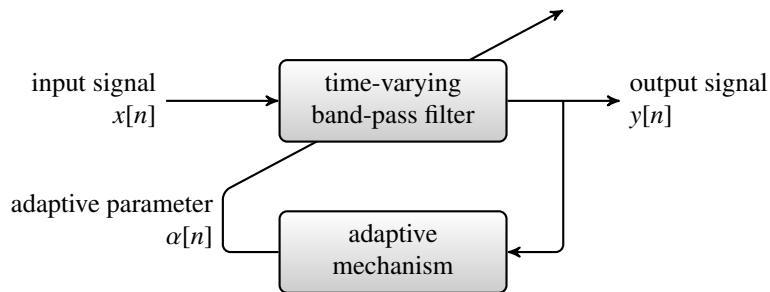


Figure 2.3: Classical structure of adaptive frequency tracking filters. In some techniques, the time-varying band-pass filter is replaced by a notch filter.

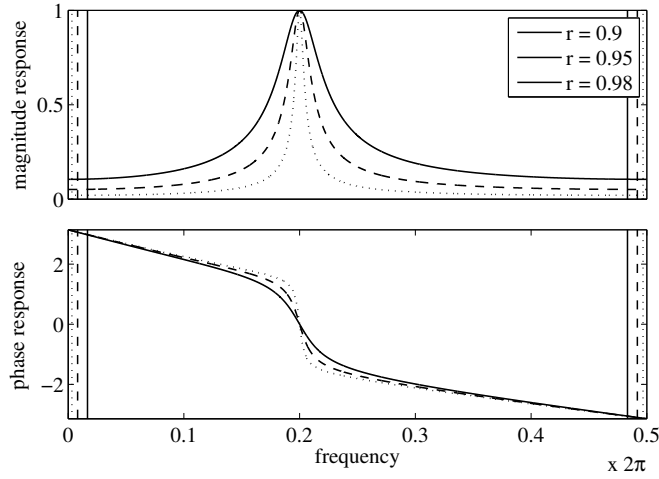


Figure 2.4: Frequency response of the band-pass filter (2.34) of the adaptive line enhancer for three values of r with the central frequency set to $\omega[n] = 0.4\pi$. The vertical lines denote the corresponding minimum and maximum possible central frequencies ω_{\min} and ω_{\max} (2.36).

This filter has unit gain and zero phase at $\omega[n]$ and therefore the extracted component is not distorted. And, due to the constraints on $\alpha[n]$, the minimum and maximum possible central frequencies are

$$\omega_{\min} = \arccos\left(\frac{2r}{1+r^2}\right) \quad \text{and} \quad \omega_{\max} = \arccos\left(\frac{-2r}{1+r^2}\right). \quad (2.36)$$

The magnitude and phase responses of the band-pass filter (2.34) are plotted in Figure 2.4 for three values of r .

Based on the filter transfer function, the output signal is

$$y[n] = \alpha[n]y[n-1] - r^2y[n-2] + \frac{1-r^2}{1+r^2}\alpha[n]x[n-1] - (1-r^2)x[n-2]. \quad (2.37)$$

And the error between input and output signals is defined as

$$e[n] = x[n] - y[n]. \quad (2.38)$$

As $y[n]$ is obtained by band-pass filtering $x[n]$, the squared error is small when the filter is centered on the oscillatory component in $x[n]$. Therefore, the goal is to minimize this quantity to estimate the instantaneous frequency. Minimizing $e^2[n]$ with a normalized gradient descent method [50] leads to the following update for the adaptive parameter:

$$\alpha[n+1] = \alpha[n] + \mu \frac{e[n]\psi[n]}{\Psi[n]}, \quad (2.39)$$

with μ ($\mu > 0$) the step-size,

$$\psi[n] = \frac{\partial y[n]}{\partial \alpha[n]} \quad (2.40)$$

and

$$\Psi[n] = \nu\Psi[n-1] + (1-\nu)\psi^2[n]. \quad (2.41)$$

The parameter ν ($0 \ll \nu < 1$) is a forgetting factor used for computing $\Psi[n]$, an estimate of the variance of $\psi[n]$. This normalization ensures that the update is scale-independent, and thus that the adaptive mechanism is robust with respect to changes in the input power. The derivative of the output signal with respect to the adaptive parameter is given by

$$\psi[n] = \alpha[n]\psi[n-1] - r^2\psi[n-2] + \frac{1-r^2}{1+r^2}x[n-1] + y[n-1]. \quad (2.42)$$

However, the performance of the algorithm is enhanced by suppressing the recursive part in the previous equation. Therefore, the expression for $\psi[n]$ becomes

$$\psi[n] = \frac{1-r^2}{1+r^2}x[n-1] + y[n-1], \quad (2.43)$$

which leads to a modified gradient that has a higher probability of pointing in the correct direction than the original one.

2.2.4 Oscillator-based Adaptive Band-pass Filter

The frequency tracking scheme presented in this section is also based on an adaptive band-pass filter. It was proposed by Liao [51], based on an earlier adaptive lattice notch filter [52]. The cost function of this algorithm is its main specificity. Indeed, it is designed to track the frequency by maximizing the oscillatory behavior of the signal under study instead of optimizing the power at the output or the difference between input and output, as done in most algorithms. Similarly to the previous technique, this scheme is composed of two parts: a time-varying band-pass filter for extracting the periodic component and an adaptive mechanism for controlling the filter. The structure is shown in Figure 2.3.

The input signal of the algorithm is defined in (2.33). The output signal, $y[n]$, is obtained by filtering the input signal with a time-varying band-pass filter whose transfer function is given by

$$H(z; n) = \frac{1-\beta}{2} \frac{1-z^{-2}}{1-\alpha[n](1+\beta)z^{-1} + \beta z^{-2}}, \quad (2.44)$$

where β ($0 \ll \beta < 1$) determines the bandwidth and $\alpha[n]$ is the adaptive parameter that controls the central frequency. This parameter is linked to the filter central frequency $\omega[n]$, which is also the frequency estimate at time n , through $\alpha[n] = \cos(\omega[n])$. The filter has zero phase and unit gain at $\omega[n]$. The influence of β on the frequency response is shown in Figure 2.5 for a central frequency set to $\omega[n] = 0.4\pi$.

The mechanism updates the filter by minimizing a cost function derived from the real discrete oscillator equation,

$$\begin{aligned} s[n] &= 2 \cos(\omega_0)s[n-1] - s[n-2] \\ &= 2\alpha_0s[n-1] - s[n-2]. \end{aligned} \quad (2.45)$$

This equation is verified for any sinusoid of the form $s[n] = A_0 \sin(\omega_0 n + \phi_0)$. From (2.45), the frequency of a pure sinusoid can be recovered with as few as three samples:

$$\omega_0 = \arccos(\alpha_0), \quad \alpha_0 = \frac{s[n] + s[n-2]}{s[n-1]}. \quad (2.46)$$

The cost function of the algorithm follows from this observation except that it takes into account interfering noise. It is defined using the filtered output signal $y[n]$:

$$J = E \left\{ (y[n] - 2\alpha[n+1]y[n-1] + y[n-2])^2 \right\}. \quad (2.47)$$

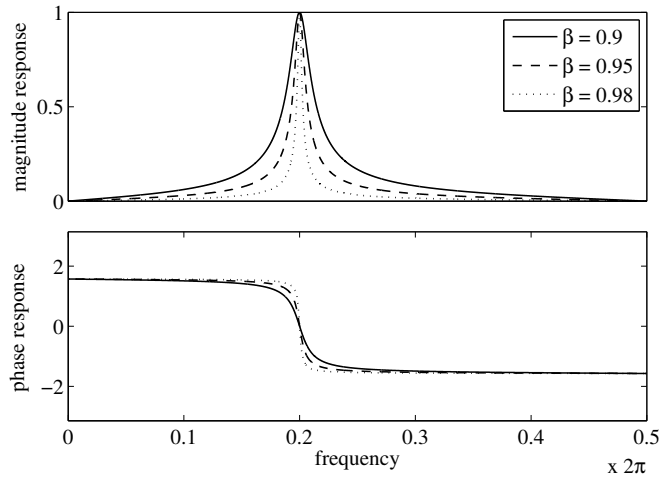


Figure 2.5: Frequency response of the oscillator-based adaptive band-pass filter (2.44) for three values of β with the central frequency set to $\omega[n] = 0.4\pi$.

Setting the derivative of this cost function with respect to $\alpha[n + 1]$ to zero, the optimal value of the adaptive parameter is given by

$$\alpha[n + 1] = \frac{\text{E}\{y[n - 1](y[n] + y[n - 2])\}}{2 \text{E}\{y^2[n - 1]\}}. \quad (2.48)$$

However, this expression is not computable in practice and therefore the expectations are replaced by exponentially weighted averages. The adaptive parameter $\alpha[n + 1]$ is updated as follows,

$$\alpha[n + 1] = \frac{Q[n]}{2P[n]} \quad (2.49)$$

with

$$Q[n] = \delta Q[n - 1] + (1 - \delta)y[n - 1](y[n] + y[n - 2]), \quad (2.50a)$$

$$P[n] = \delta P[n - 1] + (1 - \delta)y^2[n - 1], \quad (2.50b)$$

where δ ($0 \ll \delta < 1$) is a forgetting factor that controls the convergence rate. Finally, the estimate of the instantaneous frequency is computed as

$$\omega[n + 1] = \arccos(\alpha[n + 1]). \quad (2.51)$$

2.2.5 Generalized Adaptive Notch Filter

The generalized adaptive notch filter is an algorithm designed to track quasi-periodically varying complex-valued systems [53, 54]. Such systems are, in some sense, a generalization of the problem of tracking the instantaneous frequency of a single oscillation. Consequently, this technique can also be applied to the estimation of the frequency of a cisoid. In this case, it becomes an ‘ordinary’ adaptive notch filter. And, although this scheme was developed in the complex signal framework, it can still be applied to the analytic representation of a real signal obtained with the DHT, or even directly to the real signal in question.

The input signal for this method is defined as

$$x[n] = A[n]e^{j\sum_{m=1}^n \varphi[m]} + v[n], \quad n = 1, 2, \dots, \quad (2.52)$$

where the complex amplitude $A[n]$ and frequency $\varphi[n]$ are assumed to vary slowly with time, and $v[n]$ is a complex circular noise with variance σ^2 , i.e.,

$$\text{E}\{\text{Re}\{v[n]\}^2\} = \text{E}\{\text{Im}\{v[n]\}^2\} = \sigma^2/2 \quad \text{and} \quad \text{E}\{\text{Re}\{v[n]\} \text{Im}\{v[m]\}\} = 0 \quad \forall n, m. \quad (2.53)$$

The instantaneous frequency estimate $\omega[n]$ and the filtered output signal $y[n]$ are computed with the following adaptive procedure:

$$\epsilon[n] = x[n] - e^{j\omega[n]}y[n-1], \quad (2.54a)$$

$$y[n] = e^{j\omega[n]}y[n-1] + (1-\lambda)\epsilon[n], \quad (2.54b)$$

$$g[n] = \text{Im}\{\bar{\epsilon}[n]e^{j\omega[n]}y[n-1]\}, \quad (2.54c)$$

$$\omega[n+1] = \omega[n] - \eta g[n], \quad (2.54d)$$

where λ is a forgetting factor ($0 \ll \lambda < 1$), η is a step-size coefficient ($\eta > 0$) set close to zero in practice, and the upper bar denotes complex conjugation. It is worth mentioning that inserting (2.54a) into (2.54b) leads to

$$y[n] = \lambda e^{j\omega[n]}y[n-1] + (1-\lambda)x[n].$$

Therefore, the output signal $y[n]$ is obtained by filtering the input signal $x[n]$ with a time-varying band-pass filter whose transfer function is given by

$$H(z; n) = \frac{1-\lambda}{1-\lambda e^{j\omega[n]}z^{-1}}. \quad (2.55)$$

Consequently, this algorithm can be considered as an adaptive line enhancer with the structure shown in Figure 2.3. Also, this adaptive notch filter achieves the posterior Cramér-Rao bound [55] for a single noisy cisoid with frequency varying according to the random walk model under specific assumptions, namely the amplitude is constant ($A[n] = A_0$), $v[n]$ is a zero-mean complex circular Gaussian white noise, and frequency increments ($w[n] = \varphi[n] - \varphi[n-1]$) follow a zero-mean Gaussian distribution and are independent from $v[n]$. In this case, this technique is a statistically efficient procedure for tracking the instantaneous frequency.

This adaptive scheme is in fact a special case of the algorithm presented in [53] which can track K periodic components simultaneously. For multi-frequency signals, the updating procedure takes the following form:

$$\epsilon[n] = x[n] - \mathbf{1}_K^T \mathbf{\Omega}[n] y[n], \quad (2.56a)$$

$$\mathbf{P}[n] = \frac{1}{\lambda} \bar{\mathbf{\Omega}}[n] \left(\mathbf{P}[n-1] - \frac{\mathbf{P}[n-1] \mathbf{\Omega}[n] \mathbf{1}_K \mathbf{1}_K^T \bar{\mathbf{\Omega}}[n] \mathbf{P}[n-1]}{\lambda + \mathbf{1}_K^T \bar{\mathbf{\Omega}}[n] \mathbf{P}[n-1] \mathbf{\Omega}[n] \mathbf{1}_K} \right) \mathbf{\Omega}[n], \quad (2.56b)$$

$$y[n] = \mathbf{\Omega}[n] y[n-1] + \epsilon[n] \bar{\mathbf{P}}[n] \mathbf{1}_K, \quad (2.56c)$$

$$g_k[n] = \text{Im}\{\bar{\epsilon}[n] e^{j\omega_k[n]} y_k[n-1]\}, \quad k = 1, \dots, K, \quad (2.56d)$$

$$\omega_k[n+1] = \omega_k[n] - \eta g_k[n], \quad k = 1, \dots, K, \quad (2.56e)$$

with $\mathbf{1}_K = [\underbrace{1, \dots, 1}_K]^T$, $\mathbf{\Omega}[n] = \text{diag}\{e^{j\omega_1[n]}, \dots, e^{j\omega_K[n]}\}$, and $y[n] = [y_1[n], \dots, y_K[n]]^T$. The number of components K must be known.

2.2.6 Complex Adaptive Notch Filter

Another method for frequency tracking based on adaptive filters was proposed by Regalia [56]. This algorithm is based on a complex notch filter, and thus is meant to estimate the instantaneous frequency of noisy cisoids. However, as for the previous adaptive scheme, it is still applicable to a real signal or to its analytic representation. Similarly to the adaptive line enhancer, the structure of this method is composed of two parts (Figure 2.3): a time-varying filter and an adaptive mechanism.

The input signal is considered to be a cisoid embedded in noise:

$$x[n] = A_0 e^{j\omega_0 n} + v[n], \quad (2.57)$$

where A_0 and ω_0 are the complex amplitude and frequency of the cisoid, and $v[n]$ is a complex circular Gaussian noise process with variance

$$\sigma^2 = E\{|v[n]|^2\} = E\{\text{Re}\{v[n]\}^2\} + E\{\text{Im}\{v[n]\}^2\}. \quad (2.58)$$

The adaptive method is based on a first-order all-pass filter with the following transfer function,

$$C(z) = \frac{e^{j\theta} z^{-1} - \beta}{1 - \beta e^{j\theta} z^{-1}}, \quad (2.59)$$

where β ($0 \ll \beta < 1$) is the pole radius and θ is the pole angle. From this transfer function, two filters are derived, a notch filter and a band-pass filter:

$$G(z) = \frac{1}{2} (1 - C(z)) = \frac{1 + \beta}{2} \cdot \frac{1 - e^{j\theta} z^{-1}}{1 - \beta e^{j\theta} z^{-1}}, \quad (2.60a)$$

$$H(z) = \frac{1}{2} (1 + C(z)) = \frac{1 - \beta}{2} \cdot \frac{1 + e^{j\theta} z^{-1}}{1 - \beta e^{j\theta} z^{-1}}. \quad (2.60b)$$

In these transfer functions, θ determines the central frequency. The magnitude and phase responses of both notch and band-pass filters (2.60) are shown in Figure 2.6 for $\beta = 0.95$ and $\theta = 0.4\pi$. Moreover, the filters satisfy the two relations,

$$G(e^{j\omega}) + H(e^{j\omega}) = 1 \quad \text{and} \quad |G(e^{j\omega})|^2 + |H(e^{j\omega})|^2 = 1, \quad (2.61)$$

for all ω [57].

In order to track the instantaneous frequency in the input signal $x[n]$, the central frequency of the filters θ is given by the frequency estimate $\omega[n]$ at time n . Therefore, using a state-space description, the output of the notch $y[n]$ is obtained as follows,

$$u[n+1] = \beta e^{j\omega[n]} u[n] + \sqrt{1 - \beta^2} e^{j\omega[n]} x[n], \quad (2.62)$$

$$y[n] = -\frac{\sqrt{1 - \beta^2}}{2} u[n] + \frac{1 + \beta}{2} x[n]. \quad (2.63)$$

The corresponding oscillatory component can be extracted with the band-pass filter $H(z)$ (2.60b). The current estimate of the instantaneous frequency is updated with the following adaptation procedure,

$$\omega[n+1] = \omega[n] + \mu \text{Im}\{y[n] \bar{u}[n]\}, \quad (2.64)$$

where μ is the adaptation step-size. A conservative bound for μ that ensures monotonic convergence in the mean is defined as follows,

$$0 < \mu < \frac{1}{|A_0|^2} \left(\frac{1 - \beta}{1 + \beta} \right)^{3/2}. \quad (2.65)$$

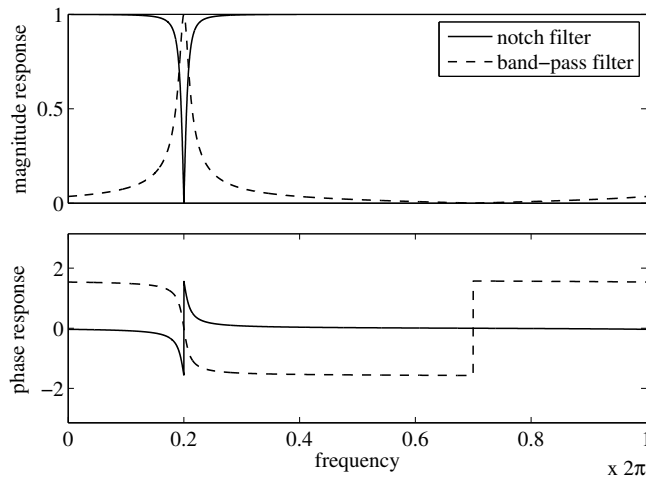


Figure 2.6: Frequency responses of the notch and band-pass filters (2.60) with pole radius and central frequency set to $\beta = 0.95$ and $\theta = 0.4\pi$ respectively.

2.2.7 Robust Modified Newton Algorithm

Unlike the four adaptive schemes presented previously, the algorithm presented in this section is not based on a time-varying filter. Instead, it relies on linear prediction for estimating the instantaneous frequency of several periodic components simultaneously. The interest of linear prediction for frequency estimation was already exploited by the classical Pisarenko's method [5] as well as the MUSIC [6] and ESPRIT [7] algorithms. However, as these techniques are based on eigenvalue decomposition, their computational load is heavy and they become inefficient in non-stationary environments. A recursive version of Pisarenko's method was proposed [58] but, despite its computational efficiency, it cannot handle multi-frequency estimation. On the other hand, highly accurate linear prediction algorithms for multiple frequency components were developed [59]. Yet, they are difficult to apply to non-stationary data due to their high computational complexity. By contrast, the robust modified Newton algorithm for adaptive frequency estimation [60] still achieves good estimation performance while offering reduced computational complexity.

The input signal for this algorithm is composed of K sinusoids embedded in white noise:

$$x[n] = s[n] + v[n] = \sum_{k=1}^K A_k \sin(\omega_k n + \phi_k) + v[n], \quad (2.66)$$

where A_k , ω_k and ϕ_k are the amplitude, frequency and phase of the k th sinusoid, and $v[n]$ is an additive white noise. Assuming that K is known, the goal is to estimate the frequencies ω_k from $y[n]$. For this purpose, one can use the fact that $s[n]$ can be perfectly predicted with an autoregressive (AR) model of order $2K$ [61]:

$$s[n] = - \sum_{k=1}^{2K} a_k s[n-k], \quad (2.67)$$

where a_k are the linear prediction coefficients. The frequencies are then equal to the phases of

the roots of the $2K$ -order polynomial

$$f(z) = \sum_{k=0}^{2K} a_k z^k \quad (2.68)$$

with $a_0 = 1$ and $a_k = a_{2K-k}$. Or, in other words,

$$\sum_{k=0}^{2K} a_k z^k = 0 \quad (2.69)$$

for $z = \exp(\pm j\omega_k)$, $k = 1, 2, \dots, K$. Based on the AR model (2.67), an error signal is defined as

$$e[n] = \sum_{k=0}^{K-1} b_k (x[n-k] + x[n-2K+k]) + b_K x[n-K], \quad (2.70)$$

where b_k are the estimates of the coefficients a_k . Unbiased estimates can be computed by minimizing the following cost function [59]:

$$J_1 = \frac{\mathbf{b}^T \mathbf{R} \mathbf{b}}{\mathbf{b}^T \mathbf{\Sigma} \mathbf{b}}, \quad (2.71)$$

with $\mathbf{\Sigma} = \text{diag}\{2, \dots, 2, 1\}$, $\mathbf{b} = [b_0, \dots, b_K]^T$, $\mathbf{x}[n] = [x[n] + x[n-2K], x[n-1] + x[n-2K+1], \dots, x[n-K+1] + x[n-K-1], x[n-K]]^T$, and $\mathbf{R} = \text{E}\{\mathbf{x}[n]\mathbf{x}^T[n]\}$. The optimal coefficient vector is given by the eigenvector corresponding to the smallest eigenvalue in the corresponding generalized eigenvalue problem:

$$\mathbf{R} \mathbf{b} = \lambda \mathbf{\Sigma} \mathbf{b}. \quad (2.72)$$

As it is difficult to use directly the cost function (2.71), an alternative cost function, defined as

$$J_2 = \frac{1}{2} (\mathbf{b}^T \mathbf{R} \mathbf{b} - \log(\mathbf{b}^T \mathbf{\Sigma} \mathbf{b})), \quad (2.73)$$

is considered instead. In fact, it can be proven that it is equivalent to minimize either of the two cost functions (2.71) and (2.73) [60].

Now, a modified Newton algorithm is derived to update the coefficient vector $\mathbf{b}[k]$ using the gradient $\mathbf{g}(\mathbf{b})$ and Hessian matrix $\mathbf{H}(\mathbf{b})$ of J_2 :

$$\mathbf{b}[k] = \mathbf{b}[k-1] - \mathbf{H}^{-1}(\mathbf{b}) \mathbf{g}(\mathbf{b}) \Big|_{\mathbf{b}=\mathbf{b}[k-1]} \quad (2.74)$$

with

$$\mathbf{g}(\mathbf{b}) = \mathbf{R} \mathbf{b} - (\mathbf{b}^T \mathbf{\Sigma} \mathbf{b})^{-1} \mathbf{\Sigma} \mathbf{b}, \quad (2.75)$$

$$\mathbf{H}(\mathbf{b}) = \mathbf{R} - \frac{\mathbf{\Sigma}}{\mathbf{b}^T \mathbf{\Sigma} \mathbf{b}} - \frac{2 \mathbf{\Sigma} \mathbf{b} \mathbf{b}^T \mathbf{\Sigma}}{(\mathbf{b}^T \mathbf{\Sigma} \mathbf{b})^2}. \quad (2.76)$$

In order to simplify the computation, the second term of the expression for the Hessian matrix is dropped, and $\mathbf{H}(\mathbf{b})$ is approximated as

$$\mathbf{H}(\mathbf{b}) \approx \mathbf{R} - \frac{2 \mathbf{\Sigma} \mathbf{b} \mathbf{b}^T \mathbf{\Sigma}}{(\mathbf{b}^T \mathbf{\Sigma} \mathbf{b})^2}. \quad (2.77)$$

Applying the matrix inversion lemma [62] to this approximation yields

$$\mathbf{H}^{-1}(\mathbf{b}) = \mathbf{R}^{-1} - \frac{2 \mathbf{R}^{-1} \mathbf{\Sigma} \mathbf{b} \mathbf{b}^T \mathbf{\Sigma} \mathbf{R}^{-1}}{(\mathbf{b}^T \mathbf{\Sigma} \mathbf{b})^2 + 2 \mathbf{b}^T \mathbf{\Sigma} \mathbf{R}^{-1} \mathbf{\Sigma} \mathbf{b}}. \quad (2.78)$$

A modified Newton algorithm is obtained by inserting the previous expression for the inverse of the Hessian matrix into the update for the coefficient vector (2.74):

$$\mathbf{b}[k] = \frac{3\mathbf{R}^{-1}\boldsymbol{\Sigma}\mathbf{b}(\mathbf{b}^T\boldsymbol{\Sigma}\mathbf{b})}{(\mathbf{b}^T\boldsymbol{\Sigma}\mathbf{b})^2 + 2\mathbf{b}^T\boldsymbol{\Sigma}\mathbf{R}^{-1}\boldsymbol{\Sigma}\mathbf{b}} \Big|_{\mathbf{b}=\mathbf{b}[k-1]}. \quad (2.79)$$

It can be shown that, although an approximate expression for the Hessian matrix (2.77) was used, this update still converges to the true coefficients [60]. However, it is not applicable in practice due to the expectation in the expression for \mathbf{R} . Thus, \mathbf{R} is replaced by an exponentially weighted sample correlation matrix defined by

$$\mathbf{R}[n] = \mu\mathbf{R}[n-1] + \mathbf{x}[n]\mathbf{x}^T[n], \quad (2.80)$$

where μ ($0 \ll \mu < 1$) is a forgetting factor. Furthermore, as the inverse of $\mathbf{R}[n]$ is used in the update, it makes more sense to compute directly $\mathbf{R}^{-1}[n]$. Therefore, letting $\mathbf{Q}[n] = \mathbf{R}^{-1}[n]$, the robust modified Newton algorithm for frequency estimation is

$$\mathbf{Q}[n] = \frac{1}{\mu} \left(\mathbf{Q}[n-1] - \frac{\mathbf{Q}[n-1]\mathbf{x}[n]\mathbf{x}^T[n]\mathbf{Q}[n-1]}{\mu + \mathbf{x}^T[n]\mathbf{Q}[n-1]\mathbf{x}[n]} \right), \quad (2.81)$$

$$\mathbf{h}[n-1] = \frac{\boldsymbol{\Sigma}\mathbf{b}[n-1]}{\mathbf{b}^T[n-1]\boldsymbol{\Sigma}\mathbf{b}[n-1]}, \quad (2.82)$$

$$\mathbf{b}[n] = \frac{3\mathbf{Q}[n]\mathbf{h}[n-1]}{1 + 2\mathbf{h}^T[n-1]\mathbf{Q}[n]\mathbf{h}[n-1]}. \quad (2.83)$$

The simplest initialization for this algorithm is to set $\mathbf{b}[0] = [1, 0, \dots, 0]$ and $\mathbf{Q}[n] = \xi\mathbf{I}$, with \mathbf{I} the identity matrix and $\xi > 0$. As this technique does not band-pass filter the input signal, it can be susceptible to noise in low SNR scenarios. It can also be mentioned that the robust modified Newton algorithm is closely related to the classical recursive least squares technique [8].

2.3 Performance Comparison

The techniques presented in the previous sections are compared in terms of frequency estimation performance. However, due to their differences, it is difficult to analyze all these algorithms on a fair basis. Indeed, depending on the considered test signal, the method achieving the best performance might change when different parameters are selected. Furthermore, the very concept of performance is rather vague as it can refer to estimation bias, estimation variance or convergence speed. For instance, in a given application the bias could be extremely important, whereas in another one the convergence speed would be the main strength of the selected algorithm. With these considerations in mind, the frequency estimation techniques were separated into two groups, the instantaneous ones and the adaptive ones, in order to compare them as fairly as possible. The first group contained the discrete Hilbert transform with phase central finite differences (DHT, Section 2.2.1) and the discrete energy separation algorithm (DESA, Section 2.2.2). On the other hand, the second group was composed of the adaptive line enhancer (ALE, Section 2.2.3), the oscillator-based adaptive band-pass filter (OSC, Section 2.2.4), the generalized adaptive notch filter (GANF, Section 2.2.5), the complex adaptive notch filter (CANF, Section 2.2.6) and the robust modified Newton algorithm for frequency estimation (RMNA, Section 2.2.7). Both groups were analyzed with Monte Carlo simulations and simple input signals.

For the first group, the input signal was a sinusoid at constant frequency 0.4π with a phase offset uniformly distributed between 0 and 2π embedded in white Gaussian noise. The DHT and DESA were applied in order to estimate the instantaneous frequency. The bias and variance of the estimates provided by the two methods were computed for SNR values ranging from 10

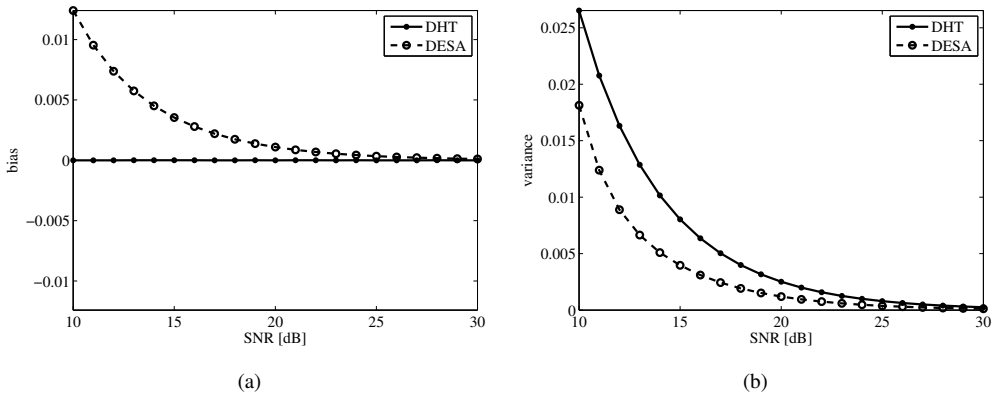


Figure 2.7: Frequency estimation bias (a) and variance (b) of the DHT and DESA for a sinusoid at 0.4π embedded in white Gaussian noise.

to 30 dB in 1 dB steps. Lower SNR values were not considered as both the DHT and DESA performed extremely poorly below 10 dB; they even failed to converge in some cases. The bias and variance were computed over the central 1000 samples of 2000 iterations in order to avoid any border effect. The two statistics averaged over 10,000 runs are plotted in Figure 2.7. The numerical simulations show that the DHT provided an unbiased frequency estimate, while the DESA estimate was clearly biased. Nevertheless, with the exception of SNR values below 15 dB, this bias remained limited. By contrast, the DESA achieved lower estimation variance than the DHT for all SNRs. To summarize, both techniques have advantages and disadvantages, and ultimately the choice should depend on the features of the signal under study.

The input signal considered for the second group was slightly different from the one for the first group. Indeed, it was also a sinusoid with a uniformly distributed phase embedded in white Gaussian noise, except that its instantaneous frequency changed abruptly from 0.2π to 0.4π . The SNR was set to 10 dB. The parameters of the ALE, OSC, GANF, CANF and RMNA were carefully selected in order to compare them as fairly as possible. First, the pole radius was set to 0.95 for the four adaptive schemes based on time-varying band-pass filters (ALE, OSC, GANF and CANF). Figure 2.8 shows that the filters have indeed very similar frequency responses, particularly around the central frequency 0.4π , once their bandwidth parameters are set accordingly. Then, the various parameters controlling the convergence rate were selected such that, once the five adaptive algorithms converged after the frequency shift, their frequency estimation variances were approximately equal. The complete set of parameters for all the adaptive techniques is summarized in Table 2.1. It is important to note that the GANF and CANF were applied to the analytic representation of the input signal obtained through the DHT. The two instantaneous methods, the DHT and DESA, were also applied in order to relate the two groups. However, the algorithms from the first group did not achieve estimation variances similar to the ones from the second group. In fact, their variances were much higher. Finally, the instantaneous frequency estimates provided by all the considered algorithms were averaged over 10,000 runs. The resulting frequency trajectories are plotted in Figure 2.9. The corresponding biases, variances and mean squared errors (MSEs) were computed over 1000 samples, 2500 iterations after the change in frequency. These statistics, averaged over 10,000 runs, are reported in Table 2.2. The results of the Monte Carlo simulations for assessing the convergence speed clearly show that both the DHT and DESA estimated very precisely the abrupt change from 0.2π to 0.4π . However, their estimation variances were much higher than for the other techniques. Furthermore, the DESA estimate

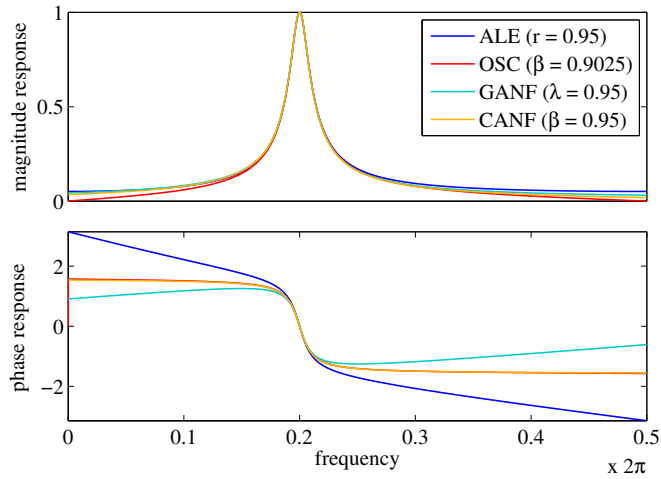


Figure 2.8: Frequency responses of time-varying filters from ALE, OSC, GANF and CANF with pole radius and central frequency set to 0.95 and 0.4π respectively.

Algorithm	Parameters
ALE	$r = 0.95, \mu = 0.006268, \nu = 0.95$
OSC	$\beta = 0.9025, \delta = 0.9328$
GANF	$\lambda = 0.95, \eta = 0.003335$
CANF	$\beta = 0.95, \mu = 0.00053$
RMNA	$K = 1, \mu = 0.995$

Table 2.1: Selected parameters for the performance comparison of the adaptive algorithms for instantaneous frequency estimation.

was clearly biased. In particular, the bias was very important before the frequency change. By contrast, all adaptive algorithms required some time to adjust. As expected, they achieved almost the same level of frequency estimation variance, and their absolute biases remained very small, nearly negligible. However, the differences in terms of convergence speed were rather important. As a rough measure of convergence speed, an algorithm was declared to have converged when its averaged frequency trajectory remained in the interval $0.4\pi \pm 0.001$. The ALE and OSC were the first to converge after 190 and 232 samples respectively. Then came the GANF with 1243 samples, the RMNA with 1306 samples, and the CANF with 1316 samples. It should be noted that the RMNA approached the true frequency much more quickly than the GANF and CANF. In fact, with the interval $0.4\pi \pm 0.005$, the following number of samples were required for convergence: 135 for the ALE, 168 for the OSC, 988 for the RMNA, 1229 for the GANF and 1271 for the CANF.

To compare convergence speeds, the GANF and CANF were applied to the analytic representation of the input signal. However, as mentioned in their respective description, they can also be applied directly to real data. It is not clear when these adaptive algorithms perform better in terms of convergence speed: when applied to the original real signal or to its analytic representation? The same procedure as before was used to answer this question, except that the GANF and CANF were applied to both the real input signal and its analytic representation computed with the DHT. Likewise, the pole radius of the time-varying filters was set to 0.95 ($\lambda = 0.95$

Algorithm	Bias	Variance	MSE
DHT	$-2.83 \cdot 10^{-6}$	$2.65 \cdot 10^{-2}$	$2.65 \cdot 10^{-2}$
DESA	$1.24 \cdot 10^{-2}$	$1.81 \cdot 10^{-2}$	$1.83 \cdot 10^{-2}$
ALE	$5.53 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$
OSC	$7.10 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$
GANF	$2.55 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$
CANF	$2.54 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$
RMNA	$-1.94 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$	$1.13 \cdot 10^{-5}$

Table 2.2: Biases, variances and mean squared errors (MSEs) of the algorithms for instantaneous frequency estimation.

for GANF and $\beta = 0.95$ for CANF), and the parameters controlling the convergence rate were set such that the two algorithms had approximately equal estimation variances. Therefore, the parameters were set as follows: $\eta = 0.005$ (GANF) and $\mu = 0.0007815$ (CANF) for the analytic representation, and $\eta = 0.01853$ (GANF) and $\mu = 0.002921$ (CANF) for the real signal. The frequency trajectories averaged over 10,000 Monte Carlo simulations are shown in Figure 2.10. Both adaptive algorithms converged approximately 250 samples faster when applied to the analytic representation than when applied to the real signal. Furthermore, the frequency estimation bias was about 16 times larger in absolute value for the real signal. In conclusion, the GANF and CANF should be applied to the analytic representation of a real signal, whenever possible.

2.4 Summary

The strengths and weaknesses of the algorithms presented in Section 2.2 are now quickly discussed. General aspects are tackled first, and then more specific features are highlighted. It should be mentioned first that none of the presented methods requires *a priori* knowledge of the considered signal. The adaptive schemes can benefit from a well-chosen initial frequency estimate, but they will eventually converge from any starting point. The instantaneous techniques, the DHT and DESA, have several features in common. Obviously, they react instantaneously to any changes in frequency, and thus they are well-suited for signals with abrupt frequency shifts. However, this reactivity comes at the cost of a very low robustness to noise. Indeed, even for relatively high SNRs the frequency estimation variance is quite important for these two methods, as illustrated with Monte Carlo simulations in Section 2.3. The bias of the DESA can also become a problem in some situations. Furthermore, when the input signal contains more than one periodic component, the DHT and DESA both fail completely and yield an irrelevant instantaneous frequency estimate (see Figure 2.1). Also, the DHT and DESA do not extract the periodic component corresponding to the estimated frequency, but instead they provide an estimate of the instantaneous amplitude (or envelope).

Similarly to the instantaneous methods, the adaptive algorithms share several features. First, they all require some time to adjust to changes in frequency, but they are not identical in terms of convergence speed. The ALE and OSC converge quickly, while the GANF, CANF and RMNA are slower. Nevertheless, they should all provide reliable estimates when applied to signals with slow frequency variations. By contrast, the two fastest schemes have slightly higher estimation biases. However, the bias for the five adaptive algorithms can be considered negligible in most cases, as it has almost no effect on the MSE (see Table 2.2). An important aspect of the techniques based on time-varying filters is that the periodic component corresponding to the frequency estimate is directly extracted, and thus it is readily available for further processing. On the other hand, the RMNA is based on linear prediction, not on adaptive filters, and does not

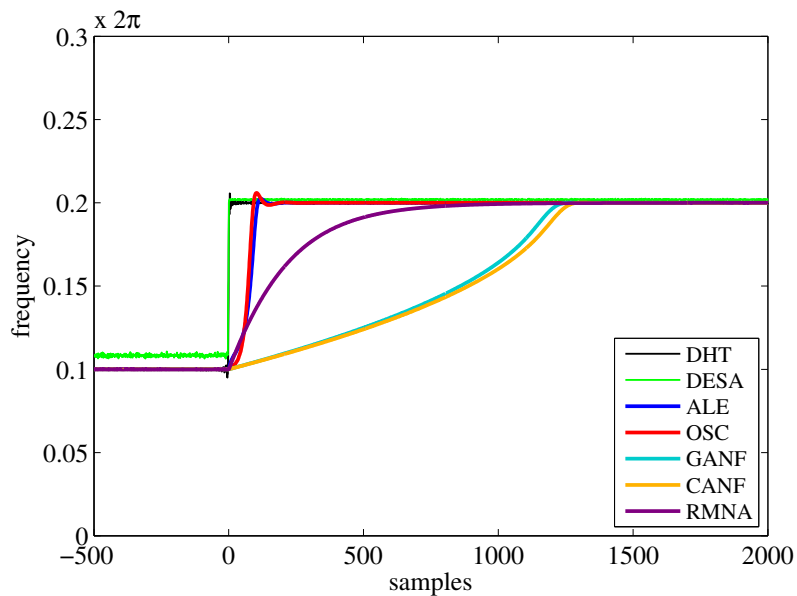


Figure 2.9: Convergence speeds of the algorithms for instantaneous frequency estimation. The gray line denotes the true frequency.

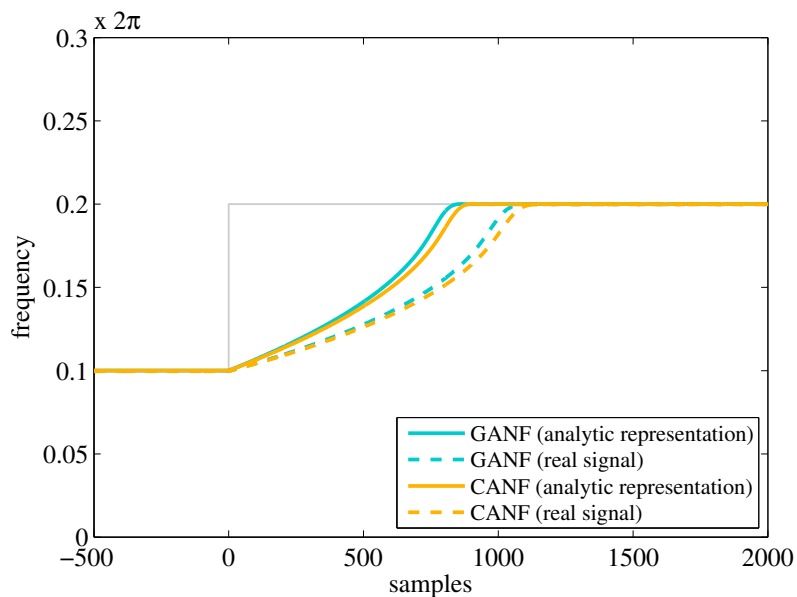


Figure 2.10: Convergence speeds of the GANF and CANF when applied to a real input signal and its analytic representation.

provide such a band-pass filtered output signal. This lack of filtering is the main cause of its susceptibility to broad-band noise compared to the other adaptive algorithms. Indeed, in the ALE, OSC, GANF and CANF, the adaptation is performed based on the signal at the output of the time-varying band-pass (or notch) filter. Consequently, a large part of the noise is discarded if the filter is sufficiently narrow. Input signals containing multiple periodic components should not be a problem for the adaptive algorithms. Indeed, the GANF can be generalized in order to take into account multiple periodic components, and the RMNA can estimate several instantaneous frequencies at the same time, as long as the number of oscillations is known. The ALE, OSC and CANF should fail gracefully; in the presence of several components, they will just track one of them (typically the most powerful one).

A final point to discuss is whether the presented techniques are applicable in real time. In fact, all the presented techniques, except the DHT, can provide an online estimate of the instantaneous frequency under certain conditions. First of all, as the DHT is typically computed with the DFT, it cannot yield a real-time frequency estimate. Nevertheless, the DHT can also be approximated with an FIR filter, and thus provide an estimate with a slight delay. This delay depends on the length of the filter, which in turn determines the accuracy of the DHT approximation. Regarding the DESA, it is not real-time *per se* as the frequency estimate at time n depends on the signal at time $n + 2$. However, in most situations, this 2-sample delay should not cause any problem. It is also important to note that, although the GANF and CANF perform better when applied to the analytic representation of a real signal, they can still provide a slightly degraded frequency estimate in real time. The main features of all the presented algorithms are summarized in Table 2.3.

DHT	DESA	ALE	OSC	GANF	CANF	RMNA
+ Instantaneous reactivity	+ Instantaneous reactivity	+ Fast convergence	+ Fast convergence	+ Component extraction	+ Component extraction	+ No <i>a priori</i> knowledge
+ Amplitude extraction	+ Amplitude extraction	+ Component extraction	+ Component extraction	+ High noise robustness	+ High noise robustness	+ Applicable in real time
+ No <i>a priori</i> knowledge	+ No <i>a priori</i> knowledge	+ High noise robustness	+ High noise robustness	+ No <i>a priori</i> knowledge	+ No <i>a priori</i> knowledge	+ Multi-frequency estimation
– Very low noise robustness	+ Applicable in real time	+ No <i>a priori</i> knowledge	+ No <i>a priori</i> knowledge	+ Applicable in real time	+ Applicable in real time	– Slow convergence
– No component extraction	– Very low noise robustness	+ Applicable in real time	+ Applicable in real time	+ Multi-frequency extension	– Slow convergence	– Low noise robustness
– Not applicable in real time	– No component extraction			– Slow convergence		– No component extraction

Table 2.3: Features of the algorithms for instantaneous frequency estimation presented in Section 2.2

Single Frequency Tracker

3

This chapter presents the single frequency tracker (SFT), an adaptive frequency tracking algorithm that was developed during this thesis. A short introduction is first given in Section 3.1. Then, the algorithm is described in detail in Section 3.2. Its performance is thoroughly analyzed in Section 3.3, and its tracking capabilities are illustrated with examples in Section 3.4. The SFT is also compared to other existing methods for instantaneous frequency estimation in Section 3.5. Finally, its strengths and weaknesses are discussed in Section 3.6. This frequency tracking technique was the subject of two publications [63, 64].

3.1 Introduction

The adaptive frequency tracking algorithm presented in this chapter is a generalization to the complex-valued signal framework of a real-valued scheme proposed by Liao [51]. The original scheme is described in Section 2.2.4. The basic idea of maximizing the oscillatory behavior at the output instead of the power is preserved but adapted to the complex discrete oscillator equation. Moreover, switching to complex-valued signals simplifies several aspects of the computations. For instance, filters are shorter with this approach. It is clear that in practice most signals are real-valued. But, with the discrete Hilbert transform (DHT), one obtains the analytic representation [34], whose real part is the original signal itself. Therefore, reverting back to real-valued signals is always possible.

3.2 Algorithm

Similarly to the original scheme, the SFT is composed of two complementary parts: a time-varying band-pass filter for extracting the oscillatory component, and an adaptive mechanism that controls the filter. The complete structure is summarized in Figure 3.1.

The input signal is defined as follows,

$$x[n] = c[n] + v[n] = A_0 e^{j\omega_0 n} + v[n], \quad (3.1)$$

where A_0 and ω_0 are the complex amplitude and frequency of the cisoid, and $v[n]$ is an additive centered complex noise. The output signal, $y[n]$, is obtained by filtering the input signal with a time-varying band-pass filter with the following transfer function:

$$G(z; n) = \frac{1 - \beta}{1 - \beta e^{j\omega[n]} z^{-1}}, \quad (3.2)$$

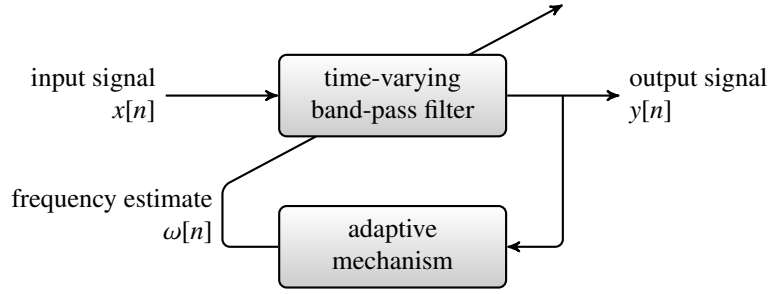
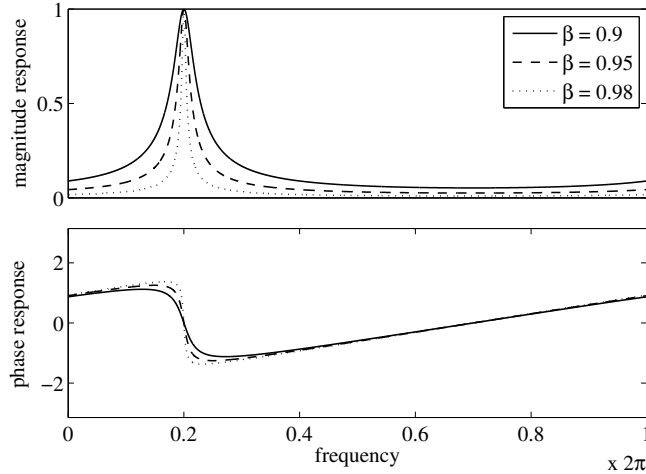


Figure 3.1: Structure of the SFT.

Figure 3.2: Response of the band-pass filter (3.2) for three values of β with the central frequency set to $\omega[n] = 0.4\pi$.

where β and $\omega[n]$ are the modulus and argument of the pole. In other words, β ($0 \ll \beta < 1$) determines the bandwidth and $\omega[n]$, which is the estimate of the instantaneous frequency at time n , controls the central frequency. The frequency response of the filter is shown in Figure 3.2 for three values of β when the central frequency is set to $\omega[n] = 0.4\pi$. This filter has zero phase and unit gain at $\omega[n]$ ensuring that the extracted oscillatory component is not distorted. Its 3-dB bandwidth is given by the following expression,

$$\Delta\omega_{3\text{ dB}} = 2 \arccos\left(\frac{1 + \beta^2}{2\beta} (1 - 10^{3/10}) + 10^{3/10}\right). \quad (3.3)$$

The adaptive mechanism is based on a cost function derived from the complex discrete oscillator equation which is defined as

$$c[n] = e^{j\omega_0} c[n-1]. \quad (3.4)$$

This equation is verified by any cisoid of the form $c[n] = A_0 e^{j\omega_0 n}$. Moreover, the frequency of such a noiseless cisoid can be estimated with only two samples using the following relation,

$$\omega_0 = \arg\{c[n]\bar{c}[n-1]\}, \quad (3.5)$$

where the upper bar denotes the complex conjugate. In a noisy environment, this simple approach is not directly applicable. Therefore, a cost function derived from (3.4) is used instead. It is defined for the output of the band-pass filter (3.2),

$$J = E \left\{ |y[n] - e^{j\omega[n+1]}y[n-1]|^2 \right\}. \quad (3.6)$$

In order to minimize this cost function, it is differentiated with respect to $\omega[n+1]$,

$$\begin{aligned} \frac{\partial J}{\partial \omega[n+1]} &= E \left\{ j e^{-j\omega[n+1]} y[n] \bar{y}[n-1] - j e^{j\omega[n+1]} \bar{y}[n] y[n-1] \right\} \\ &= j e^{-j\omega[n+1]} E \{ y[n] \bar{y}[n-1] \} - j e^{j\omega[n+1]} E \{ \bar{y}[n] y[n-1] \} \\ &= 2 \operatorname{Re} \left\{ j e^{-j\omega[n+1]} E \{ y[n] \bar{y}[n-1] \} \right\} \\ &= -2 \operatorname{Im} \left\{ e^{-j\omega[n+1]} E \{ y[n] \bar{y}[n-1] \} \right\}. \end{aligned} \quad (3.7)$$

Setting this derivative to zero yields the optimal value for $\omega[n+1]$:

$$\omega[n+1] = \arg \left\{ E \{ y[n] \bar{y}[n-1] \} \right\}. \quad (3.8)$$

However, this expression is not applicable in practice. Thus, the expectation is replaced by an exponentially weighted average [65], and the adaptive mechanism becomes

$$Q[n] = \delta Q[n-1] + (1 - \delta) y[n] \bar{y}[n-1], \quad (3.9)$$

$$\omega[n+1] = \arg \{ Q[n] \}, \quad (3.10)$$

where $Q[n]$ is an internal variable and δ ($0 \ll \delta < 1$) is a forgetting factor that determines the convergence rate.

3.3 Performance Analysis

Three aspects of the performance of the SFT are analyzed now. First of all, the behavior of the cost function (3.6) is studied in detail in Section 3.3.1. Then, the frequency estimation bias and variance of the adaptive scheme are investigated in depth in Sections 3.3.2 and 3.3.3. Whenever possible, theoretical results are compared to numerical simulations. It should be noted that only the final results are presented here. Nonetheless, the performance of the SFT is analyzed in detail in Appendix B where all derivations can be found.

All the performance analyses are carried out for a generic cisoid embedded in noise since the SFT is based on this specific input signal. This generic noisy cisoid is defined as

$$y[n] = A_0 e^{j\omega_0 n} + v[n], \quad (3.11)$$

where A_0 is the complex amplitude, ω_0 is the frequency, and $v[n]$ is an additive wide-sense stationary zero-mean noise process [8]. The other statistical properties of the noise, such as its autocorrelation $R_{vv}[k] = E\{v[n]\bar{v}[n-k]\}$, are not defined yet in order to remain as general as possible. Indeed, the cost function, the bias and the variance are studied first for the generic noisy cisoid, and then the results are refined with the features of four specific input signals. The first of the considered signals is a simple cisoid with additive white noise. The band-pass filter (3.2) is not applied in this case in order to investigate the behavior of the adaptive mechanism only. Then, the second considered signal is the analytic representation, computed with the DHT, of a real noisy sinusoid. Indeed, real data are much more common in practical situations. In addition, like for the first signal, no filtering is performed. The last two input signals are the

same as the first two except that they are band-pass filtered before the instantaneous frequency is estimated. Also, since in some part of the analyses the central frequency is fixed, the transfer function of the band-pass filter is rewritten here in a slightly modified form:

$$G(e^{j\omega}) = \frac{1 - \beta}{1 - \beta e^{j\omega_c} e^{-j\omega}}, \quad (3.12)$$

where ω_c is the central frequency. To summarize, the four considered input signals and their respective noise autocorrelation are defined as follows.

Noisy cisoid A cisoid embedded in additive zero-mean complex white noise:

$$y_1[n] = A_0 e^{j\omega_0 n} + v_1[n], \quad (3.13)$$

where the noise variance is σ^2 . Furthermore, the real and imaginary parts of the noise are assumed to be independent and to have equal variances. The noise autocorrelation is

$$R_{v_1 v_1}[k] = \sigma^2 \delta[k], \quad (3.14)$$

with σ^2 the variance and $\delta[k]$ the Kronecker delta (A.5).

Noisy sinusoid The analytic representation of a sinusoid embedded in additive zero-mean real white noise:

$$y_2[n] = A_0 e^{j\omega_0 n} + v_2[n], \quad (3.15)$$

where $v_2[n]$ is the analytic representation of a real white noise process with variance σ^2 . The autocorrelation of this representation is given by

$$R_{v_2 v_2}[k] = 2\sigma^2 (\delta[k] + jh[k]), \quad (3.16)$$

with $h[k]$ the Hilbert filter (2.12).

Filtered noisy cisoid A cisoid embedded in additive zero-mean complex white noise filtered with the band-pass filter (3.12):

$$y_3[n] = \frac{1 - \beta}{1 - \beta e^{-j(\omega_0 - \omega_c)}} A_0 e^{j\omega_0 n} + v_3[n], \quad (3.17)$$

with $v_3[n] = g[n] * v_1[n]$ and $g[n]$ the impulse response of (3.12). The notation $*$ in the previous expression denotes the convolution. In this case, the autocorrelation of the noise is defined as

$$R_{v_3 v_3}[k] = \sigma^2 \frac{1 - \beta}{1 + \beta} \beta^{|k|} e^{j\omega_c k}. \quad (3.18)$$

Filtered noisy sinusoid The analytic representation of a sinusoid embedded in additive zero-mean real white noise filtered with the band-pass filter (3.12):

$$y_4[n] = \frac{1 - \beta}{1 - \beta e^{-j(\omega_0 - \omega_c)}} A_0 e^{j\omega_0 n} + v_4[n], \quad (3.19)$$

with $v_4[n] = g[n] * v_2[n]$. The autocorrelation of the filtered analytic representation of the real white noise is

$$R_{v_4 v_4}[k] = 2\sigma^2 \frac{1 - \beta}{1 + \beta} e^{j\omega_c k} \left(\beta^{|k|} - \frac{j}{\pi} \left(\beta^k \ln \left(\frac{1 + \beta e^{j\omega_c}}{1 - \beta e^{j\omega_c}} \right) - \beta^{-k} \ln \left(\frac{1 + \beta e^{-j\omega_c}}{1 - \beta e^{-j\omega_c}} \right) \right) + \frac{2j}{\pi} b[k] \right), \quad (3.20)$$

with

$$b[k] = \begin{cases} 0 & \text{for } k = -1, 0, 1, \\ \beta^k \sum_{l=0}^{\lfloor k/2 \rfloor - 1} \frac{(\beta^{-1} e^{-j\omega_c})^{2l+1}}{2l+1} - \beta^{-k} \sum_{l=0}^{\lfloor k/2 \rfloor - 1} \frac{(\beta e^{-j\omega_c})^{2l+1}}{2l+1} & \text{for } k > 1, \\ \bar{b}[-k] & \text{for } k < -1. \end{cases}$$

Parameter	Default	Range
SNR	10 dB	from 0 to 20 dB in 1 dB steps
ω_0	0.4π	from 0.04π to 0.96π in 0.04π steps
ω_c	$\omega[n]$	from 0.04π to 0.96π in 0.04π steps
β	0.95	from 0.5 to 0.95 in 0.05 steps
δ	0.95	from 0.5 to 0.95 in 0.05 steps

Table 3.1: Default values and ranges for the parameters used in the performance analysis of the SFT.

Therefore, the four considered input signals are indeed special cases of the generic noisy cisoid (3.11). Furthermore, the additive noise is white only for the first signal. In the other ones, the DHT and the band-pass filter (3.12) introduce long-term correlations in the noise processes.

As both the bias and variance were investigated with Monte Carlo simulations, the analysis procedure is briefly presented. Two different versions of the SFT were applied: a simplified scheme without the adaptive band-pass filter for the first two signals, and the complete algorithm for the last two. Furthermore, two distinct types of input signals were generated. For the noisy cisoid and filtered noisy cisoid cases, the inputs were cisoids with a uniformly distributed random phase offset embedded in complex circular Gaussian white noise [66]. On the other hand, for the noisy sinusoid and filtered noisy sinusoid cases, sinusoids with uniformly distributed phase and additive real Gaussian white noise were used. The bias and variance were computed over the last 1000 samples of 2000 iterations and averaged over 10,000 runs. These two quantities were analyzed with respect to several parameters: the SNR, the cisoid or sinusoid frequency ω_0 , the central frequency ω_c of the filter (3.12), the bandwidth parameter β , and the forgetting factor δ . The numerical simulations were performed by varying one of the parameters in a given range while the other ones were fixed to their default value. These default values and ranges are summarized in Table 3.1. The central frequency ω_c can be considered as a special case since its default value is the current frequency estimate $\omega[n]$. This means that the central frequency of the band-pass was adjusted to the current estimate at all time, as described in the adaptive algorithm. By contrast, when ω_c was varied across the possible frequencies, the frequency estimation was performed with a slightly modified SFT wherein the adaptation of the filter was blocked. Thus, although the estimate was still updated, the central frequency was kept fixed. It should also be mentioned that the amplitude A_0 of the cisoids and sinusoids was set to one in all cases.

3.3.1 Cost Function Analysis

A detailed analysis of the cost function (3.6) of the SFT is required in order to properly understand the behavior of the algorithm in different situations. In particular, it is important to detect if it is biased under certain conditions. The oscillator-based cost function, repeated here for convenience, is defined as

$$J(\omega) = \text{E} \left\{ \left| y[n] - e^{j\omega} y[n-1] \right|^2 \right\}. \quad (3.21)$$

This expression is first derived for the generic cisoid by inserting (3.11) into (3.21), which becomes

$$J(\omega) = 4|A_0|^2 \sin^2 \left(\frac{\omega - \omega_0}{2} \right) + 2 \left(R_{vv}[0] - \text{Re} \{ e^{-j\omega} R_{vv}[1] \} \right). \quad (3.22)$$

where $R_{vv}[k] = \text{E}\{v[n]\bar{v}[n-k]\}$ is the noise autocorrelation, as defined previously. The bias only depends on the second term, and therefore on the noise, as the first term is minimized for $\omega = \omega_0$. It is now easy to analyze the cost function for the four considered input signals since it suffices to insert the corresponding noise autocorrelation into (3.22).

Noisy cisoid

With $R_{v_1 v_1}[k]$ (3.14), the cost function (3.22) becomes

$$J(\omega) = 4|A_0|^2 \sin^2\left(\frac{\omega - \omega_0}{2}\right) + 2\sigma^2. \quad (3.23)$$

Therefore, minimizing this expression yields an unbiased frequency estimate, as illustrated in Figure 3.3 for different noise variances.

Noisy sinusoid

Plugging the autocorrelation $R_{v_2 v_2}[k]$ (3.16) into (3.22) leads to

$$J(\omega) = 4|A_0|^2 \sin^2\left(\frac{\omega - \omega_0}{2}\right) + 4\sigma^2 \left(1 - \frac{2}{\pi} \sin(\omega)\right). \quad (3.24)$$

For this input signal, the cost function is biased towards $\omega = 0.5\pi$ as the second term is minimized for this frequency. Figure 3.4 shows the bias in this case for different noise levels. In fact, the DHT, which introduces long-term correlations in the analytic representation of the noise, is causing the bias.

Filtered noisy cisoid

Inserting $R_{v_3 v_3}[k]$ in (3.22) yields the following cost function:

$$J(\omega) = \frac{(1 - \beta)^2}{1 - 2\beta \cos(\omega_0 - \omega_c) + \beta^2} \cdot 4|A_0|^2 \sin^2\left(\frac{\omega - \omega_0}{2}\right) + 2\sigma^2 \frac{1 - \beta}{1 + \beta} (1 - \beta \cos(\omega - \omega_c)). \quad (3.25)$$

As such, minimizing this expression with respect to ω leads to an estimate biased towards the central frequency ω_c of the band-pass filter for this input signal. Nevertheless, when the central and cisoid frequencies coincide ($\omega_c = \omega_0$), the cost function reduces to

$$J(\omega) = 4|A_0|^2 \sin^2\left(\frac{\omega - \omega_0}{2}\right) + 2\sigma^2 \frac{1 - \beta}{1 + \beta} (1 - \beta \cos(\omega - \omega_0)), \quad (3.26)$$

and the resulting frequency estimate is unbiased. The effects of the noise variance and central frequency are shown in Figure 3.5. Accordingly, when the cisoid and central frequency are in close agreement, the bias is very limited. It even vanishes when they match. It should also be noted that this case corresponds to the complete SFT with band-pass filtering compared to the two previous input signals. Consequently, as long as the frequency estimate provided by the SFT is close to the true cisoid frequency, the bias should remain negligible.

Filtered noisy sinusoid

Similar to the first three input signals, replacing $R_{v_v}[k]$ in (3.22) by $R_{v_4 v_4}[k]$ (3.20) leads to

$$\begin{aligned} J(\omega) = & \frac{(1 - \beta)^2}{1 - 2\beta \cos(\omega_0 - \omega_c) + \beta^2} \cdot 4|A_0|^2 \sin^2\left(\frac{\omega - \omega_0}{2}\right) \\ & + 4\sigma^2 \frac{1 - \beta}{1 + \beta} \cdot \left(1 - \beta \cos(\omega - \omega_c) + \frac{j}{\pi} \ln\left(\frac{1 - \beta^2 - 2j\beta \sin(\omega_c)}{1 - \beta^2 + 2j\beta \sin(\omega_c)}\right)\right) \\ & + \text{Im} \left\{ \frac{e^{-j(\omega - \omega_c)}}{\pi\beta} \left(\ln\left(\frac{1 + \beta e^{-j\omega_c}}{1 - \beta e^{-j\omega_c}}\right) - \beta^2 \ln\left(\frac{1 + \beta e^{j\omega_c}}{1 - \beta e^{j\omega_c}}\right) \right) \right\}. \end{aligned} \quad (3.27)$$

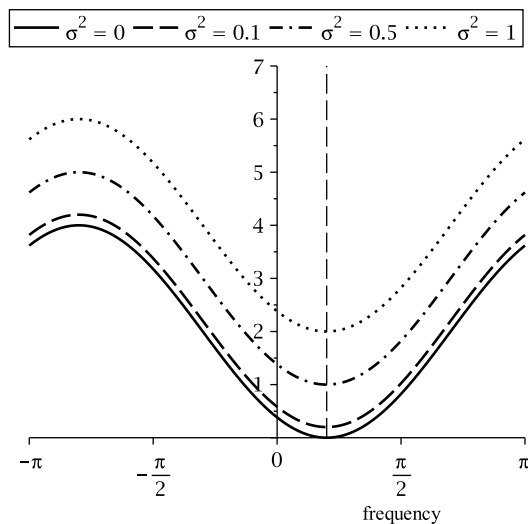


Figure 3.3: Oscillator-based cost function for a noisy cisoid. The cost function is plotted for four different noise variances. The amplitude and frequency are set to $A_0 = 1$ and $\omega_0 = 0.2\pi$. The vertical dashed line denotes the true frequency of the cisoid.

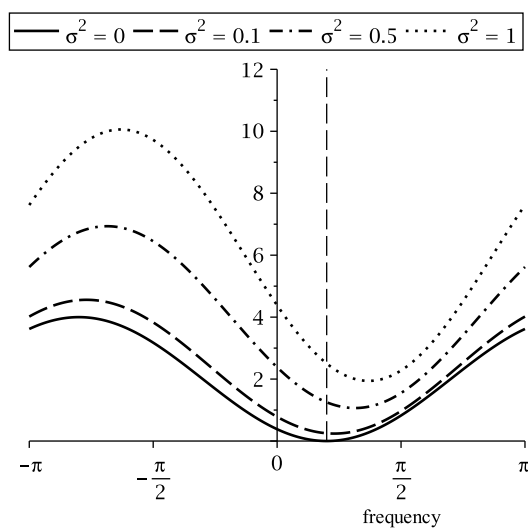


Figure 3.4: Oscillator-based cost function for the analytic representation of a noisy sinusoid. The cost function is plotted for four different noise variances. The amplitude and frequency are set to $A_0 = 1$ and $\omega_0 = 0.2\pi$. The vertical dashed line denotes the true frequency of the sinusoid.

A biased frequency estimate is obtained when minimizing this cost function. Indeed, the bias is caused by two separate factors in this case: the DHT and the band-pass filter (3.12). Nonetheless, the part of the bias due to the filtering operation is eliminated when the cisoid and central frequencies are identical, and the expression for the cost function simplifies into

$$\begin{aligned}
J(\omega) = & 4|A_0|^2 \sin^2\left(\frac{\omega - \omega_0}{2}\right) \\
& + 4\sigma^2 \frac{1 - \beta}{1 + \beta} \cdot \left(1 - \beta \cos(\omega - \omega_0) + \frac{j}{\pi} \ln\left(\frac{1 - \beta^2 - 2j\beta \sin(\omega_0)}{1 - \beta^2 + 2j\beta \sin(\omega_0)}\right)\right) \\
& + \text{Im} \left\{ \frac{e^{-j(\omega - \omega_0)}}{\pi\beta} \left(\ln\left(\frac{1 + \beta e^{-j\omega_0}}{1 - \beta e^{-j\omega_0}}\right) - \beta^2 \ln\left(\frac{1 + \beta e^{j\omega_0}}{1 - \beta e^{j\omega_0}}\right) \right) \right\}.
\end{aligned} \tag{3.28}$$

The general cost function for the input signal (3.27) is plotted for different combinations of noise variances and central frequencies in Figure 3.6. To summarize, when the central frequency of the filter matches the cisoid frequency and the bandwidth is sufficiently narrow, the bias caused by the DHT becomes negligible.

3.3.2 Bias Analysis

As discussed in the analysis of the cost function (3.22), the SFT is biased under certain conditions. In particular, the estimation bias is not equal to zero whenever the DHT is used to compute the analytic representation of a real signal or the band-pass filter (3.12) is misaligned with respect to the frequency of the periodic component under study. However, the relations between the bias magnitude, the signal features and the parameters of the SFT still need to be investigated in more detail. Therefore, a theoretical value is derived for the bias, which is then refined for the four considered input signals. In fact, only an approximation is calculated since the adaptive mechanism of the SFT computes the instantaneous frequency by taking the argument of the autocorrelation of the signal $y[n]$ at lag one, $R_{yy}[1]$, estimated with an exponentially weighted average $Q[n]$ (3.9),

$$\omega[n + 1] = \arg\{Q[n]\}. \tag{3.10}$$

Thus, the first step to derive the approximate bias is to compute the expected value of $Q[n]$, $\tilde{Q} = E\{Q[n]\}$, for the generic noisy cisoid (3.11):

$$\tilde{Q} = |A_0|^2 e^{j\omega_0} + R_{vv}[1]. \tag{3.29}$$

Based on this expression, the approximate expected value of $\omega[n]$ is defined as

$$\tilde{\omega} = \arg\{\tilde{Q}\} = \arg\{|A_0|^2 e^{j\omega_0} + R_{vv}[1]\}, \tag{3.30}$$

and the approximate bias is given by

$$\text{Bias}\{\tilde{\omega}\} = \arg\{|A_0|^2 e^{j\omega_0} + R_{vv}[1]\} - \omega_0, \tag{3.31}$$

where ω_0 is the true frequency. Therefore, if $R_{vv}[1]$ is zero or its phase is equal to ω_0 , the frequency estimate provided by the SFT is unbiased. This theoretical expression is compared to Monte Carlo simulations for the four input signals. In these simulations, the parameters of the signals and the SFT are varied according to Table 3.1.

Noisy cisoid

Analogously to the cost function analysis, the noise autocorrelation $R_{v_1 v_1}[1]$ (3.14) for this input signal is inserted into the generic expression of the bias (3.31):

$$\text{Bias}\{\tilde{\omega}\} = \arg\{|A_0|^2 e^{j\omega_0}\} - \omega_0 = 0. \tag{3.32}$$

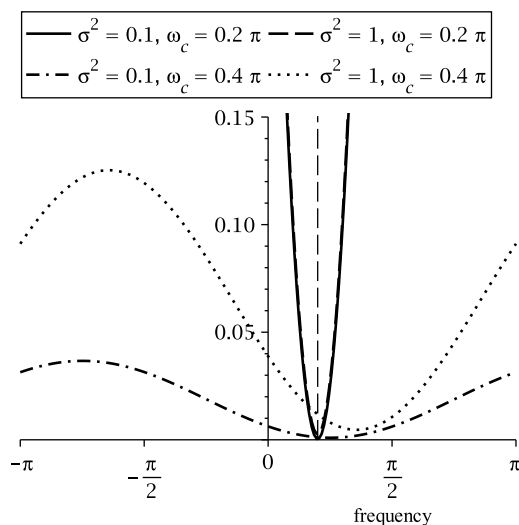


Figure 3.5: Oscillator-based cost function for a filtered noisy cisoid. The cost function is plotted for four combinations of noise variances and central frequencies. The amplitude and frequency are set to $A_0 = 1$ and $\omega_0 = 0.2\pi$, and the bandwidth parameter of the filter (3.12) is set to $\beta = 0.95$. The vertical dashed line denotes the true frequency of the cisoid.

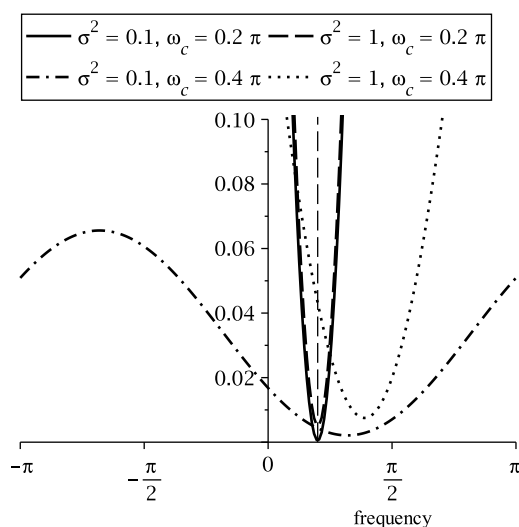


Figure 3.6: Oscillator-based cost function for the filtered analytic representation of a noisy sinusoid. The cost function is plotted for four combinations of noise variances and central frequencies. The amplitude and frequency are set to $A_0 = 1$ and $\omega_0 = 0.2\pi$, and the bandwidth parameter of the filter (3.12) is set to $\beta = 0.95$. The vertical dashed line denotes the true frequency of the sinusoid.

Thus, the frequency estimate of the simplified SFT is unbiased, which is confirmed with numerical simulations for various SNRs in Figure 3.7.

Noisy sinusoid

With $R_{v_2v_2}[1] = j4\sigma^2/\pi$, the expression for the bias (3.31) becomes

$$\text{Bias}\{\tilde{\omega}\} = \arg \left\{ |A_0|^2 e^{j\omega_0} + j \frac{4\sigma^2}{\pi} \right\} - \omega_0. \quad (3.33)$$

This theoretical value is in close agreement with the outcomes of the Monte Carlo simulations. Figure 3.8 shows this comparison for multiple SNR values and sinusoids frequencies. In particular, the frequency estimate is biased towards 0.5π (Figure 3.8b) as pointed out by the cost function analysis for this input signal. The DHT is responsible for the non-zero bias since the analytic representation of a real white noise is no longer white. Globally, the bias absolute value depended on both the SNR and the sinusoid frequency.

Filtered noisy cisoid

The autocorrelation $R_{v_3v_3}[k]$ is plugged into (3.31). This results in the following expression for the bias:

$$\text{Bias}\{\tilde{\omega}\} = \arg \left\{ \frac{(1-\beta)^2}{1-2\beta \cos(\omega_0 - \omega_c) + \beta^2} \cdot |A_0|^2 e^{j\omega_0} + \sigma^2 \frac{1-\beta}{1+\beta} \beta e^{j\omega_c} \right\} - \omega_0. \quad (3.34)$$

Therefore, the SFT yields an unbiased estimate when the central frequency of the band-pass filter (3.12) is equal to the cisoid frequency: $\omega_c = \omega_0$. The behavior of the bias is investigated with respect to the SNR and to the central frequency in Figure 3.9 for this input signal. The theoretical value matches closely the Monte Carlo simulations results in both graphs. The key point of this analysis is that, whenever the central frequency is close to the cisoid frequency, the bias of the SFT becomes negligible.

Filtered noisy sinusoid

Similarly to the previous cases, the autocorrelation $R_{v_v}[k]$ in the bias expression (3.31) is replaced with $R_{v_4v_4}[k]$ (3.20). Thus, the expected value of $Q[n]$ becomes

$$\begin{aligned} \tilde{Q} = & \frac{(1-\beta)^2}{1-2\beta \cos(\omega_0 - \omega_c) + \beta^2} \cdot |A_0|^2 e^{j\omega_0} \\ & + 2\sigma^2 \frac{1-\beta}{1+\beta} \left(\beta e^{j\omega_c} + \frac{j e^{j\omega_c}}{\pi\beta} \left(\ln \left(\frac{1+\beta e^{-j\omega_c}}{1-\beta e^{-j\omega_c}} \right) - \beta^2 \ln \left(\frac{1+\beta e^{j\omega_c}}{1-\beta e^{j\omega_c}} \right) \right) \right), \end{aligned} \quad (3.35)$$

and the estimation bias for this input signal is given by

$$\text{Bias}\{\tilde{\omega}\} = \arg\{\tilde{Q}\} - \omega_0. \quad (3.36)$$

This theoretical value is compared to the bias obtained with numerical simulations as before. More precisely, the influences of the SNR, the sinusoid frequency, the central frequency and the bandwidth parameter are investigated on both theoretical and practical biases. All results are illustrated in Figure 3.10 and, with the exception of the special case $\omega_0 = \omega_c = 0.5\pi$, the SFT is biased. The bias is caused by mismatches between central and sinusoid frequencies and by the fact that the analytic representation of a real white noise is not white. Nevertheless, as long as the band-pass filter is sufficiently narrow and centered on the sinusoid frequency, the bias of the SFT for this input signal remains limited and can even be considered negligible in most situations.

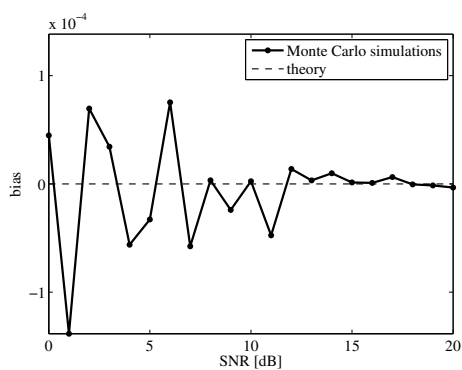
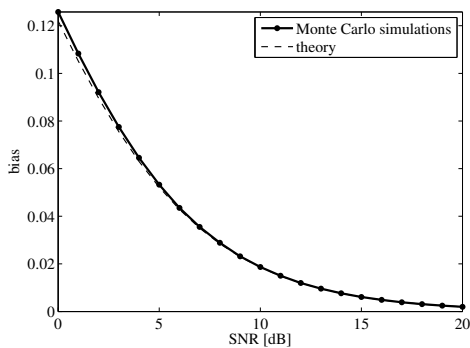
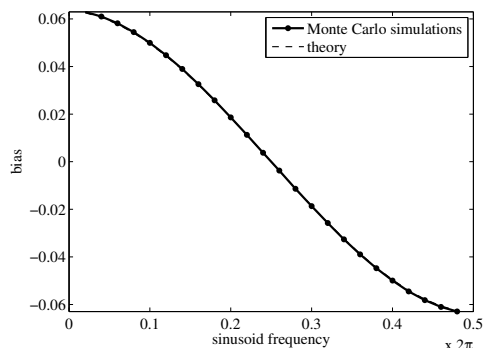


Figure 3.7: Bias of the SFT for a noisy cisoid with respect to the SNR ($A_0 = 1$, $\omega_0 = 0.4\pi$, $\delta = 0.95$).

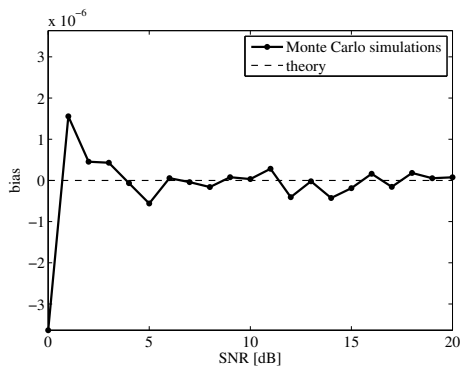


(a)

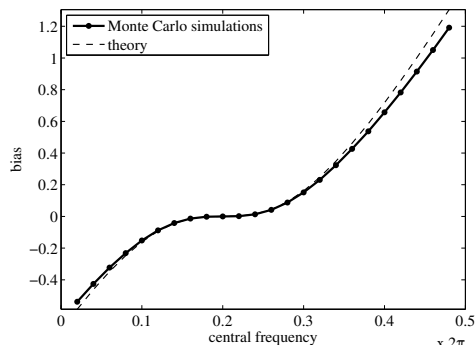


(b)

Figure 3.8: Bias of the SFT for the analytic representation of a noisy sinusoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\delta = 0.95$) with respect to (a) the SNR, and (b) to the sinusoid frequency ω_0 .



(a)



(b)

Figure 3.9: Bias of the SFT for a filtered noisy cisoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$) with respect to (a) the SNR, and (b) to the central frequency ω_c .

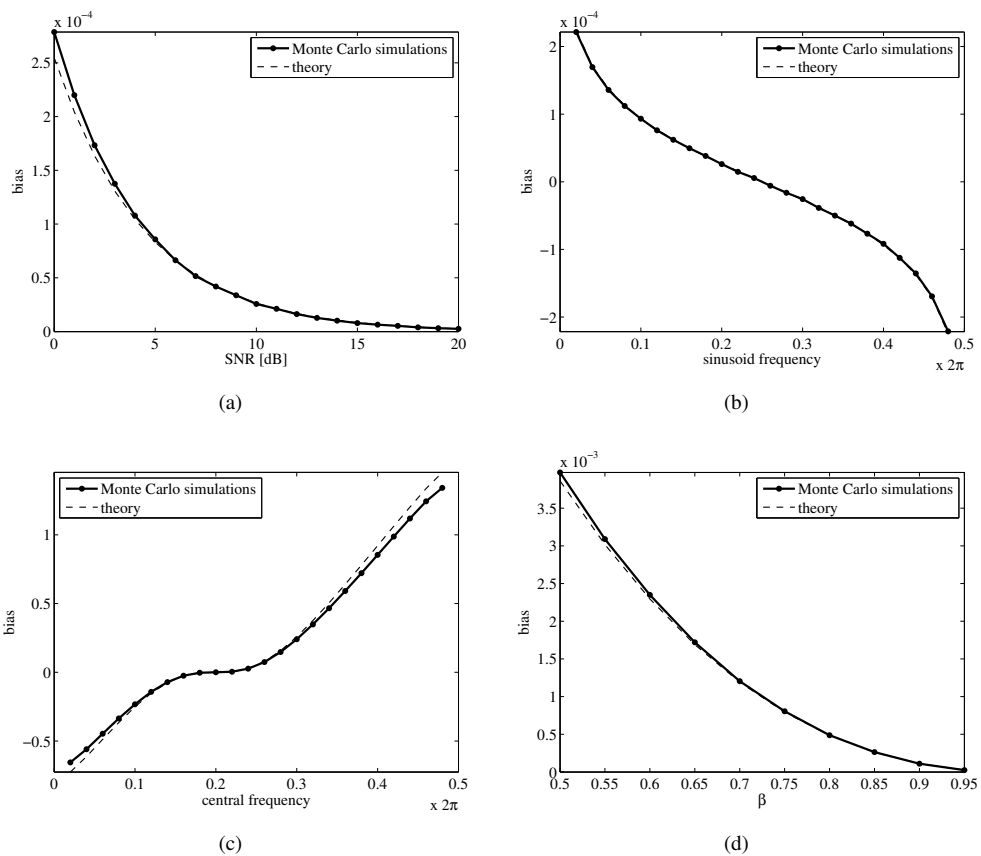


Figure 3.10: Bias of the SFT for the filtered analytic representation of a noisy sinusoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$) with respect (a) to the SNR, (b) to the sinusoid frequency ω_0 , (c) to the central frequency ω_c , and (d) to the bandwidth parameter β .

3.3.3 Variance Analysis

In the same way as in the bias analysis, the effects of the signal features and SFT parameters on the frequency estimation variance are examined for the four considered inputs. However, the theoretical calculations quickly become intractable because of the argument in the update for the frequency estimate $\omega[n]$ (3.10). Notwithstanding these difficulties, it is still possible to use Monte Carlo simulations in order to analyze the variance of the SFT. The following sections present the outcomes of these simulations for the values of the parameters given in Table 3.1. It should be noted that a theoretical expression for the variance of $Q[n]$ is still derived for the noisy cisoid, noisy sinusoid and filtered noisy cisoid cases in Section B.4.

Noisy cisoid

The results of the numerical simulations for the noisy cisoid case are shown in Figure 3.11 for various SNRs and forgetting factor δ . As expected, high SNR values lead to low frequency estimation variance. The variance is also reduced for forgetting factors close to one.

Noisy sinusoid

Like for the previous input signal, the estimation variance is investigated with respect to the SNR and to the forgetting factor. In addition to these parameters, the sinusoid frequency is also considered as it influences the bias. Figure 3.12 illustrates the outcomes of the Monte Carlo procedure. First, the same remarks about the SNR and δ as for the noisy cisoid apply. Indeed, low noise variances and forgetting factors close to one reduce the estimation variance. On the other hand, the effect of the sinusoid frequency is very limited, but a small decrease in variance around the frequency 0.5π is reminiscent of the results obtained for the bias with the same input signal.

Filtered noisy cisoid

For this input signal, the variance of $\omega[n]$ is analyzed by varying the SNR, the forgetting factor, and the central frequency and bandwidth parameter of the band-pass filter (3.12). All the results of the numerical simulations are shown in Figure 3.13. The SNR and δ have the same effects as for the two previous cases. The central frequency plays also an important role as the estimation variance increases with its distance from the cisoid frequency. Still, as long as the filter is approximately aligned with the tracked periodic component, the variance increase is almost unnoticeable. At last, the bandwidth parameter β has an influence very similar to that of the forgetting factor. Indeed, values close to one decrease the variance. This is intuitive as the closer β is to one, the narrower the band-pass filter is, thus a greater part of the noise is discarded.

Filtered noisy sinusoid

The estimation variance for the filtered noisy sinusoid is examined with respect to the same parameters as for the previous case. However, the effect of the sinusoid frequency ω_0 is also investigated with Monte Carlo simulations. All results are shown in Figure 3.14. Overall, the same comments as for the noisy sinusoid and filtered noisy cisoid can be reiterated in this case. Namely, high SNRs, forgetting factors close to one and narrow bandwidths lead to a low estimation variance. Furthermore, it also decreases when the central frequency of the filter (3.12) is in close agreement with the sinusoid frequency. In contrast with the noisy sinusoid case, ω_0 has almost no effect on the variance.

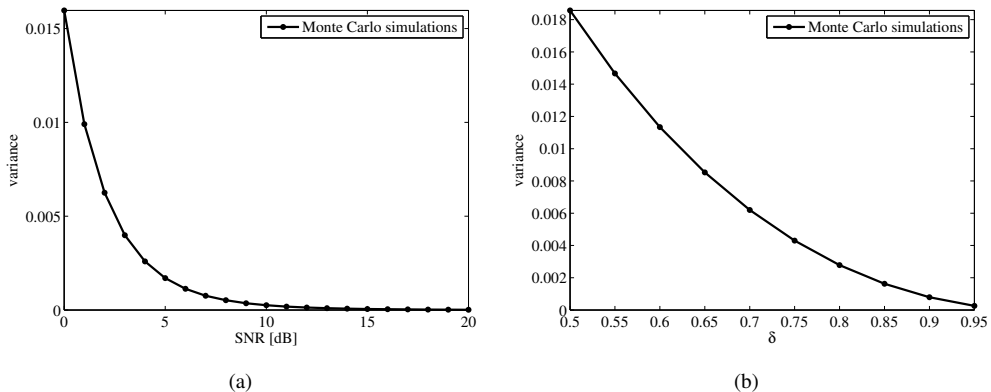


Figure 3.11: Variance of the SFT for a noisy cisoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\delta = 0.95$) with respect (a) to the SNR, and (b) to the forgetting factor δ .

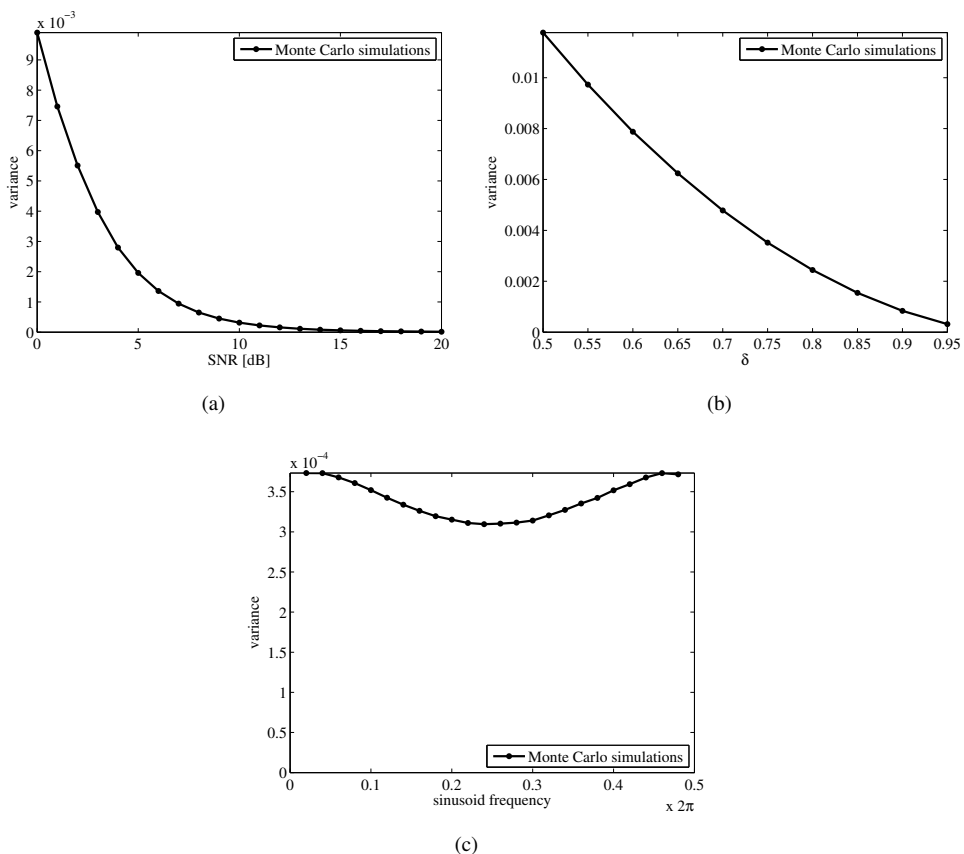


Figure 3.12: Variance of the SFT for the analytic representation of a noisy sinusoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\delta = 0.95$) with respect (a) to the SNR, (b) to the forgetting factor δ , and (c) to the sinusoid frequency ω_0 .

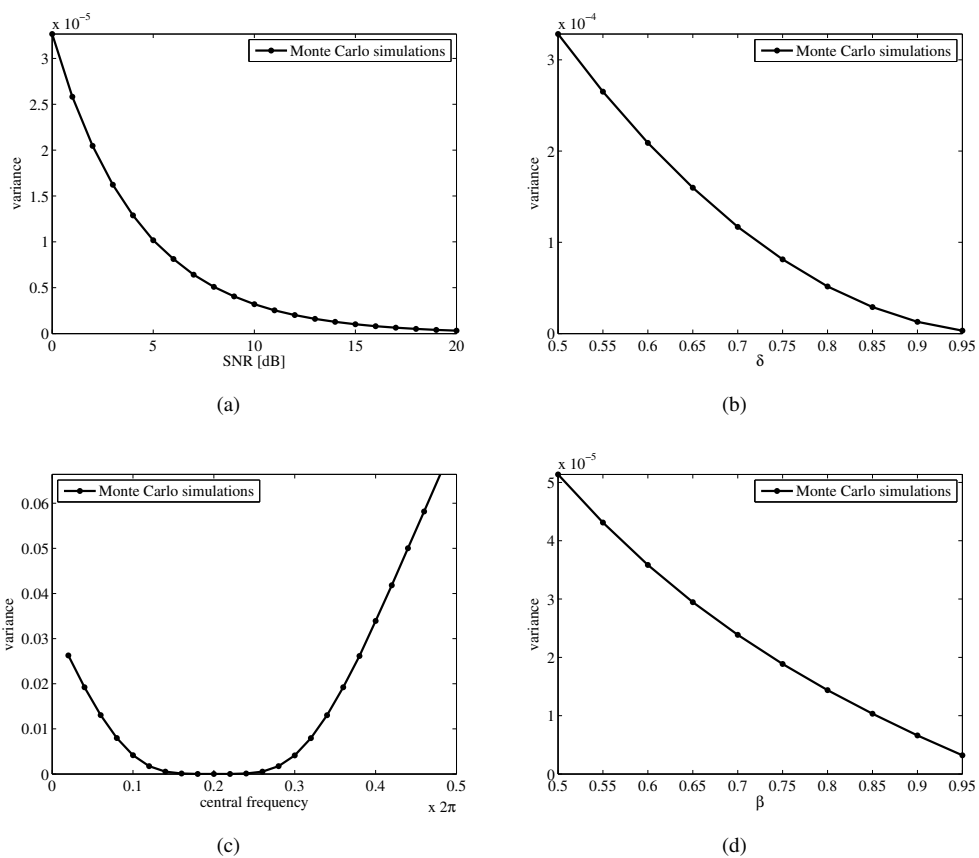


Figure 3.13: Variance of the SFT for a filtered noisy cisoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$) with respect (a) to the SNR, (b) to the forgetting factor δ , (c) to the central frequency ω_c , and (d) to the bandwidth parameter β .

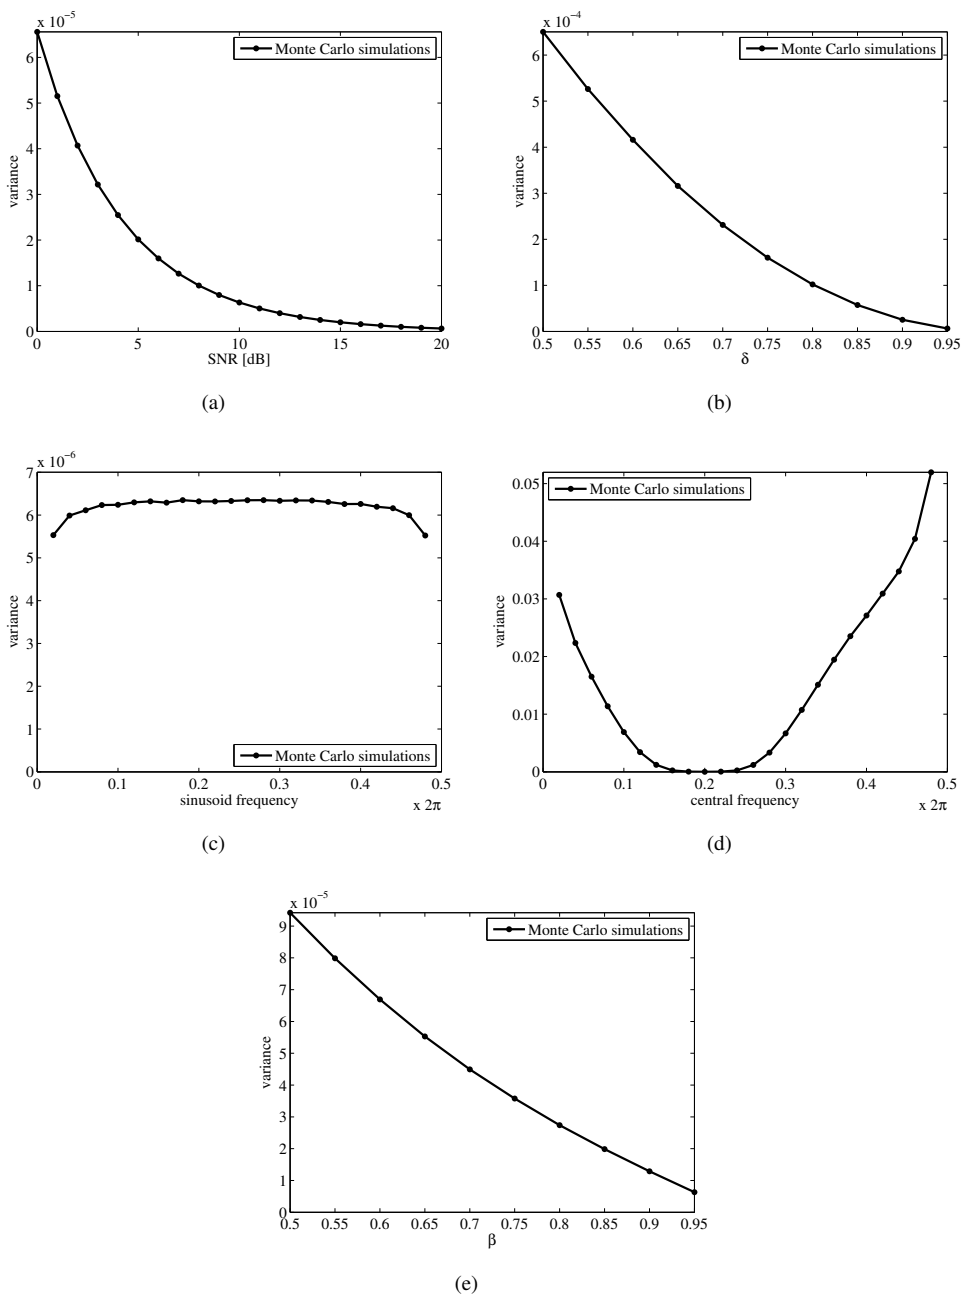


Figure 3.14: Variance of the SFT for the filtered analytic representation of a noisy sinusoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$) with respect (a) to the SNR, (b) to the forgetting factor δ , (c) to the sinusoid frequency ω_0 , (d) to the central frequency ω_c , and (e) to the bandwidth parameter β .

3.4 Frequency Tracking Examples

The behavior of the SFT is now illustrated with three examples. In the first one, the input signal is a cisoid embedded in complex circular Gaussian white noise (SNR = 5 dB). The cisoid frequency changes abruptly from 0.2π to 0.4π after 200 samples. The bandwidth parameter of the SFT is set to $\beta = 0.95$. Figure 3.15 shows the instantaneous frequency estimate provided by the SFT for three values of the forgetting factor ($\delta = 0.9, 0.95, 0.97$). In addition, the real part of the input signal is plotted in Figure 3.16 alongside the real parts of the oscillations extracted by the SFT for these three forgetting factors. The influence of δ is clearly visible with this example. A low value of this parameter leads to fast tracking, but at the cost of increased estimation variance. On the contrary, when it is close to one the tracking is slower and the estimation is more accurate.

The second example illustrates the tracking performance of the SFT for a linear chirp. The input signal is thus a cisoid whose frequency changes linearly from 0.2π to 0.6π in 500 samples with additive complex circular Gaussian white noise. The following parameters are selected for this example: SNR = 5 dB, $\beta = 0.95$, $\delta = 0.95$. The instantaneous frequency estimated by the SFT is shown in Figure 3.17 with the true value.

In the last example, the differences in tracking capabilities of the SFT are tested when it is applied to a complex signal and to the analytic representation of a real signal. This aspect of the SFT is investigated with a frequency-modulated (FM) oscillation embedded in noise. Thus, the input signal is a cisoid whose frequency oscillated between 0.3π and 0.5π with additive complex circular Gaussian white noise for the complex case, and it was the analytic representation of a noisy FM sinusoid for the real case. The following parameters are selected for both cases, SNR = 5 dB, $\beta = 0.95$, $\delta = 0.95$. The instantaneous frequencies estimated by the SFT are plotted in Figure 3.18 for the two cases, and they are almost indistinguishable. However, the tracking seems to be a little bit slower for the analytic representation of a real signal, but it is barely noticeable.

3.5 Comparisons with Existing Tracking Algorithms

As the performance of the SFT was discussed at length in Section 3.3, it is now compared to other existing frequency tracking algorithms with Monte Carlo simulations. More precisely, the SFT is confronted in terms of convergence speed to a selection of the techniques presented in Section 2.2: the adaptive line enhancer (ALE, Section 2.2.3), the oscillator-based adaptive band-pass filter (OSC, Section 2.2.4), the generalized adaptive notch filter (GANF, Section 2.2.5), and the complex adaptive notch filter (CANF, Section 2.2.6). Although the frequency estimates computed with the DHT (Section 2.2.1) and the discrete energy separation algorithm (DESA, Section 2.2.3) react instantaneously to frequency changes, their extreme susceptibility to noise requires too many constraints to obtain fair comparisons. Furthermore, the robust modified Newton algorithm for frequency estimation (RMNA, Section 2.2.7) is also left outside of these comparisons since it lacks a filtering operation for reducing the noise. And overall, it makes more sense to compare the SFT to similar techniques based on adaptive filters. Nevertheless, it is still possible to relate the DHT, DESA, RMNA and SFT by paralleling the following results with the ones presented in Section 2.3. The various frequency tracking algorithms were compared with two sets of signals: periodic components with an abrupt frequency change and oscillations whose instantaneous frequency evolves according to a random walk model. Both sets were investigated within the real- and complex-valued frameworks as three of the considered methods (SFT, GANF and CANF) are designed to process complex-valued data while the other two (ALE and OSC) require real-valued inputs. Therefore, when working with real-valued data, the SFT, GANF and CANF were applied to the analytic representations of the input signals. On the other hand, the ALE and OSC used the real parts of complex-valued signals.

The signals in the first set were defined as sinusoids or cisoids with an abrupt frequency

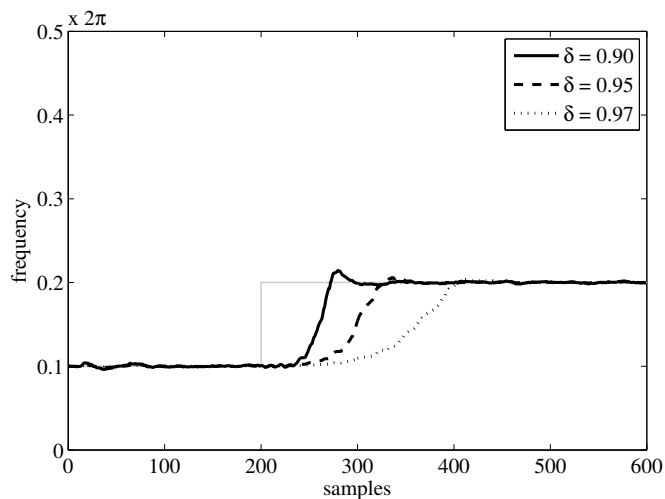


Figure 3.15: Instantaneous frequency estimated by the SFT for a noisy cisoid with an abrupt frequency shift ($\text{SNR} = 5 \text{ dB}$, $\beta = 0.95$). The estimate is shown for three values of the forgetting factor δ , and the gray line denotes the true frequency.

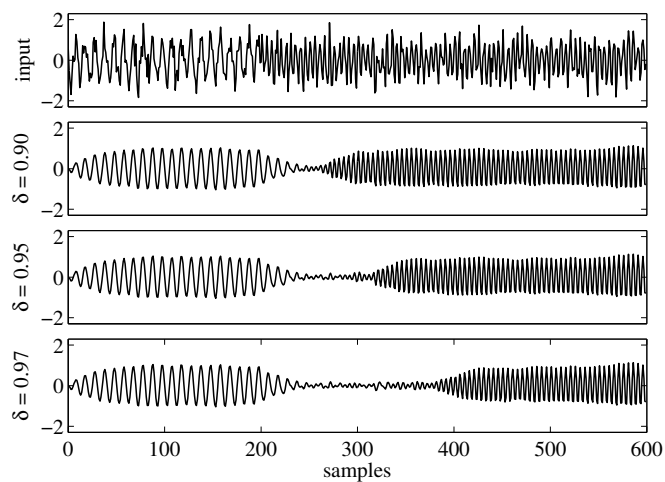


Figure 3.16: Input signal and oscillations extracted by the SFT for a noisy cisoid with an abrupt frequency shift ($\text{SNR} = 5 \text{ dB}$, $\beta = 0.95$). The oscillations were extracted with three different values of the forgetting factor δ .

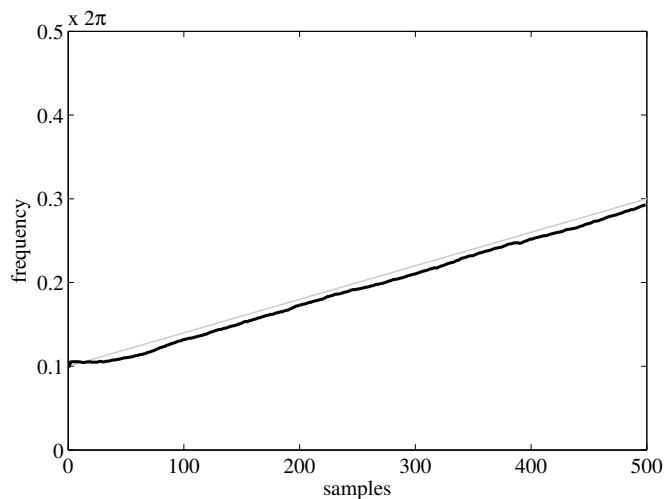


Figure 3.17: Instantaneous frequency estimated by the SFT for a linear chirp from 0.2π to 0.3π embedded in complex Gaussian white noise (SNR = 5 dB, $\beta = 0.95$, $\delta = 0.95$). The gray line denotes the true frequency.

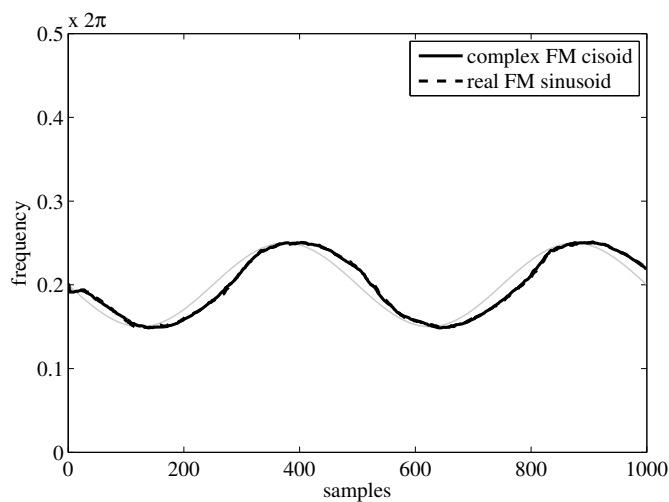


Figure 3.18: Instantaneous frequency estimated by the SFT for real and complex FM oscillations between 0.3π and 0.5π embedded in Gaussian white noise (SNR = 5 dB, $\beta = 0.95$, $\delta = 0.95$). The gray line denotes the true frequency.

Algorithm	Parameters for real-valued signals	Parameters for complex-valued signals
SFT	$\beta = 0.95, \delta = 0.95$	$\beta = 0.95, \delta = 0.95$
ALE	$r = 0.95, \mu = 0.004685, \nu = 0.95$	$r = 0.95, \mu = 0.003323, \nu = 0.95$
OSC	$\beta = 0.9025, \delta = 0.9495$	$\beta = 0.9025, \delta = 0.9641$
GANF	$\lambda = 0.95, \eta = 0.00249$	$\lambda = 0.95, \eta = 0.002495$
CANF	$\beta = 0.95, \mu = 0.000399$	$\beta = 0.95, \mu = 0.0004$

Table 3.2: Selected parameters for the convergence speed comparisons of the adaptive frequency tracking algorithms.

Algorithm	Statistics for real-valued signals			Statistics for complex-valued signals		
	Bias	Variance	MSE	Bias	Variance	MSE
SFT	$8.33 \cdot 10^{-5}$	$2.00 \cdot 10^{-5}$	$2.00 \cdot 10^{-5}$	$-1.63 \cdot 10^{-7}$	$1.02 \cdot 10^{-5}$	$1.02 \cdot 10^{-5}$
ALE	$1.61 \cdot 10^{-4}$	$2.00 \cdot 10^{-5}$	$2.01 \cdot 10^{-5}$	$1.48 \cdot 10^{-4}$	$1.02 \cdot 10^{-5}$	$1.02 \cdot 10^{-5}$
OSC	$2.37 \cdot 10^{-4}$	$2.00 \cdot 10^{-5}$	$2.01 \cdot 10^{-5}$	$2.44 \cdot 10^{-4}$	$1.02 \cdot 10^{-5}$	$1.02 \cdot 10^{-5}$
GANF	$8.57 \cdot 10^{-5}$	$2.00 \cdot 10^{-5}$	$2.00 \cdot 10^{-5}$	$-3.68 \cdot 10^{-7}$	$1.02 \cdot 10^{-5}$	$1.02 \cdot 10^{-5}$
CANF	$8.56 \cdot 10^{-5}$	$2.00 \cdot 10^{-5}$	$2.01 \cdot 10^{-5}$	$-3.52 \cdot 10^{-7}$	$1.02 \cdot 10^{-5}$	$1.02 \cdot 10^{-5}$

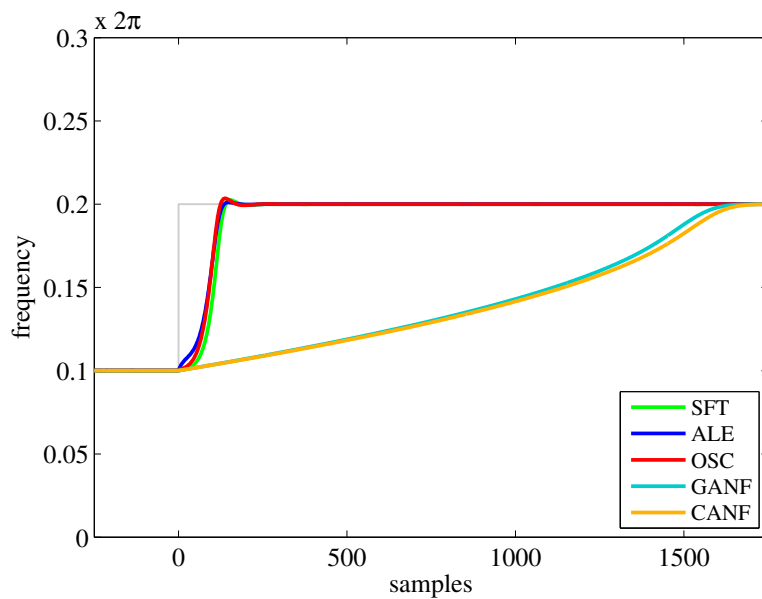
Table 3.3: Biases, variances and mean squared errors (MSEs) of the frequency tracking algorithms for the convergence speed comparisons with real- and complex-valued signals.

change from 0.2π to 0.4π and a random phase offset uniformly distributed between 0 and 2π embedded in additive white noise. In the real-valued case, the noise followed a Gaussian distribution, whereas, in the complex-valued case, it was a complex circular Gaussian white noise. The SNR was set to 5 dB in both cases. In order to obtain fair comparisons of the frequency tracking algorithms, their parameters were carefully selected. First, the pole radius of the time-varying filter of each adaptive scheme was set to 0.95. Then, the forgetting factors and step-sizes controlling the convergence rate were set such that all algorithms achieved approximately the same level of estimation variance after convergence. Table 3.2 documents the selected parameters for each technique and each case. The frequency trajectories averaged over 10,000 runs are plotted in Figure 3.19 for real- and complex-valued input signals. The biases, variances and mean squared errors (MSEs) for all algorithms, which were computed over 1000 samples after convergence and averaged over 10,000 runs as well, are reported in Table 3.3. There is not much difference between the outcomes of the numerical simulations for real- and complex-valued signals. Indeed, in both cases, the SFT, ALE and OSC converged quickly after the frequency change while the GANF and CANF were slower. The difference in terms of convergence speed between fast and slow algorithms is roughly 1500 samples, which can be important in some applications. It can also be noted that, as the SFT is designed to operate on complex-valued data, it was the first to converge for complex-valued signals while it came third for real-valued signals. The SFT achieved the smallest absolute bias in both cases, but the GANF and CANF were close. On the other hand, the bias for the two real schemes was more important.

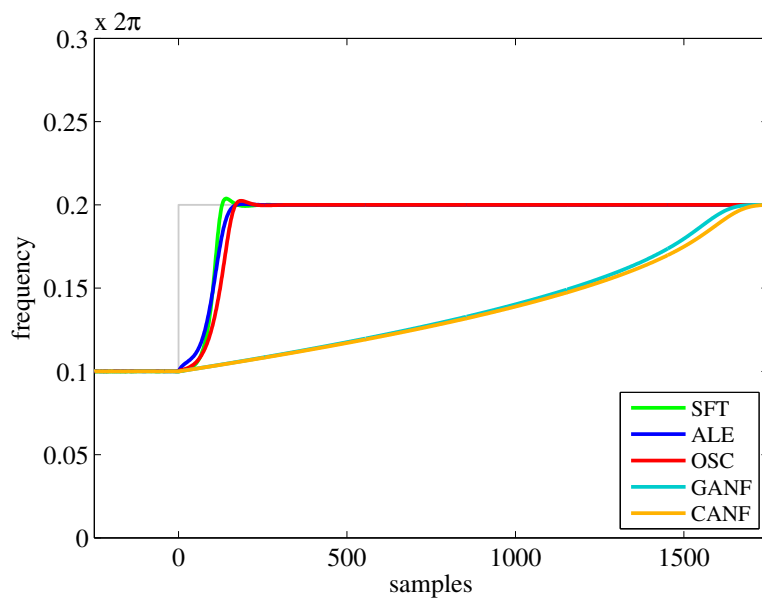
The signals in the second set were also sinusoids for the real case or cisoids for the complex case with uniformly distributed phase embedded in additive white noise. Furthermore, the noise was defined as for the first set. On the other hand, the instantaneous frequency $\omega[n]$ of the sinusoids and cisoids fluctuated according to a random walk model:

$$\omega[n] = \omega[n-1] + w[n], \quad (3.37)$$

where the frequency increments $w[n]$ were zero-mean Gaussian distributed random variables



(a) Real-valued signals.



(b) Complex-valued signals.

Figure 3.19: Convergence speeds of the frequency tracking algorithms for (a) real- and (b) complex-valued signals. The gray lines denote the true frequency.

with standard deviation $\sigma_w = 0.005$. The initial frequency was set to 0.5π . For this set of signals, the performance was evaluated by computing the MSE between the true frequency and the estimate of each adaptive scheme over the last 1000 samples of 2000 iterations for SNR values taken from 0 to 20 dB in 5 dB steps. These MSEs were averaged over 10,000 simulations. As before, the aim was to compare the five frequency tracking algorithms as fairly as possible, and thus the various parameters were scrupulously selected. Like for the signals with an abrupt frequency change, the pole radius of all adaptive filters was set to 0.95. The parameters controlling the convergence rate were chosen such that the SFT, ALE, OSC, GANF and CANF had approximately equal estimation variances when tracking the frequency of an oscillation at 0.5π embedded in noise. These parameters are given in Table 3.4 for all considered techniques and SNRs. It was confirmed with Monte Carlo simulations that the variances of all algorithms measured over the last 1000 samples of 2000 iterations and averaged over 10,000 runs were indeed approximately the same for oscillations at 0.5π . Figure 3.20 shows the MSE of each adaptive scheme when tracking a noisy sinusoid or cisoid with random walk frequency fluctuations. In these numerical simulations, the SFT performed the best for both real- and complex-valued signals. Yet, all the considered algorithms were close in terms of MSE, particularly in the real case. It can also be noted that, in the real case as well, the ALE and OSC, which are designed to process real-valued signals, performed better than the GANF and CANF. And, as expected, it was the opposite in the complex case.

3.6 Discussion

The SFT, presented in this chapter, is an adaptive algorithm for extracting a periodic component buried in noise and its instantaneous frequency. Its main specificity, compared to similar schemes based on adaptive filters, is its cost function based on the complex discrete oscillator equation (3.4). By means of this cost function, the SFT is designed to maximize the oscillatory behavior at the output, instead of, for instance, maximizing the output power or minimizing the error between input and output signals. Due to this feature and its update mechanism, the algorithm is scale-independent. And thus, it can even extract an oscillation with low amplitude compared to the noise level. Another relevant aspect of the SFT is its time-varying band-pass filter (3.2). Indeed, since it has unit gain and zero phase at the central frequency, the filtered periodic component is not distorted. Depending on the type of information that needs to be extracted, this can be of the utmost importance. For example, when quantifying the synchronization between two oscillations, it is crucial that the component extraction does not introduce any phase lag. The influence of each parameter on the performance of the presented technique was also thoroughly analyzed with both theoretical calculations and Monte Carlo simulations. The key findings are that the SFT is unbiased for a cisoid embedded in white noise and mostly unbiased for a noisy sinusoid. As the tracking algorithm is applied to the analytic representation of a real signal, the DHT introduces long-term correlations in the noise process which in turn causes a limited bias. However, as long as the filter is well-adjusted and its bandwidth is sufficiently narrow, the bias can be considered negligible in most situations. And, when compared to other existing frequency tracking algorithms based on adaptive filters, the SFT achieved excellent performance. Indeed, during the investigation of the reactivity to sharp frequency changes, the presented algorithm was the fastest for complex-valued signals and the third, but extremely close to the best one, for real-valued signals. Furthermore, when the instantaneous frequency of the components evolved according to a random walk model, the SFT yielded the lowest MSE for each tested SNR in both real and complex cases.

Some limitations of the presented approach are also worth mentioning. First, as illustrated in the abrupt frequency change example (Figure 3.15), the update of the SFT is not immediate. The algorithm needs some time to adapt its current frequency estimate and to adjust the position

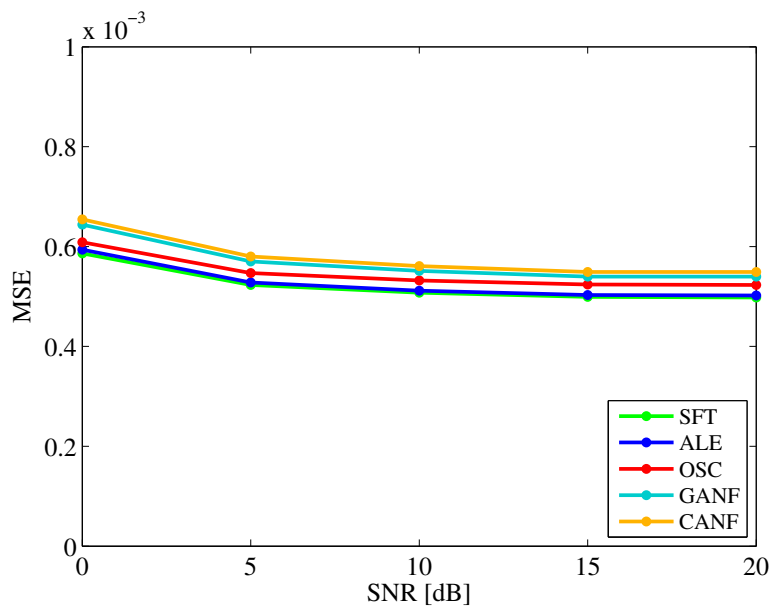
Algorithm	Fixed parameters	SNR = 0 dB	SNR = 5 dB	SNR = 10 dB	SNR = 15 dB	SNR = 20 dB
SFT	$\beta = 0.95$	$\delta = 0.95$	$\delta = 0.95$	$\delta = 0.95$	$\delta = 0.95$	$\delta = 0.95$
ALE	$r = 0.95, \nu = 0.95$	$\mu = 0.004941$	$\mu = 0.00493$	$\mu = 0.004927$	$\mu = 0.004926$	$\mu = 0.004925$
OSC	$\beta = 0.9025$	$\delta = 0.94914$	$\delta = 0.94934$	$\delta = 0.94939$	$\delta = 0.94941$	$\delta = 0.94942$
GANF	$\lambda = 0.95$	$\eta = 0.002471$	$\eta = 0.002489$	$\eta = 0.002497$	$\eta = 0.002499$	$\eta = 0.002499$
CANF	$\beta = 0.95$	$\mu = 0.0003958$	$\mu = 0.0003987$	$\mu = 0.0004$	$\mu = 0.0004004$	$\mu = 0.0004004$

(a) Parameters for real-valued signals.

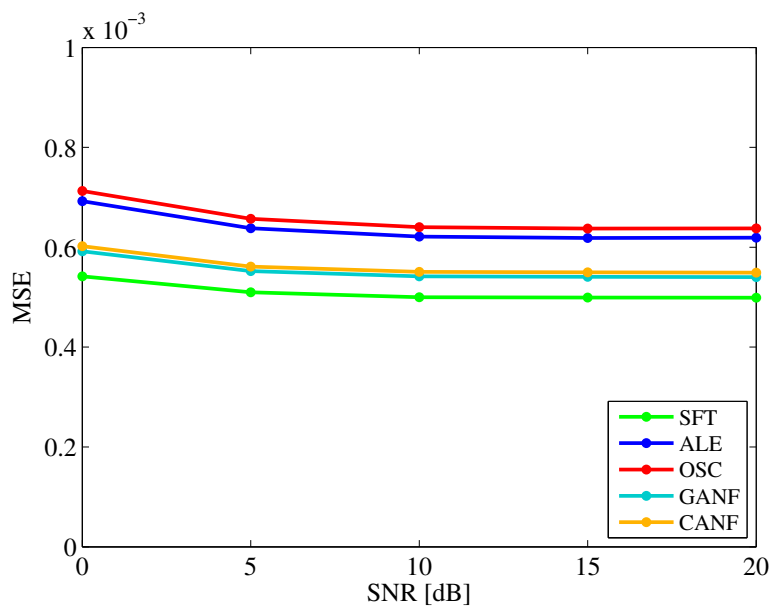
Algorithm	Fixed parameters	SNR = 0 dB	SNR = 5 dB	SNR = 10 dB	SNR = 15 dB	SNR = 20 dB
SFT	$\beta = 0.95$	$\delta = 0.95$	$\delta = 0.95$	$\delta = 0.95$	$\delta = 0.95$	$\delta = 0.95$
ALE	$r = 0.95, \nu = 0.95$	$\mu = 0.003465$	$\mu = 0.003489$	$\mu = 0.003501$	$\mu = 0.003504$	$\mu = 0.003503$
OSC	$\beta = 0.9025$	$\delta = 0.96419$	$\delta = 0.96402$	$\delta = 0.96393$	$\delta = 0.9639$	$\delta = 0.96391$
GANF	$\lambda = 0.95$	$\eta = 0.002486$	$\eta = 0.002495$	$\eta = 0.002498$	$\eta = 0.002499$	$\eta = 0.0025$
CANF	$\beta = 0.95$	$\mu = 0.0003981$	$\mu = 0.0003999$	$\mu = 0.0004003$	$\mu = 0.0004006$	$\mu = 0.0004007$

(b) Parameters for complex-valued signals.

Table 3.4: Selected parameters for the mean squared error comparisons of the adaptive frequency tracking algorithms for random walk frequency fluctuations for (a) real- and (b) complex-valued signals.



(a) Real-valued signals.



(b) Complex-valued signals.

Figure 3.20: Mean squared errors (MSEs) of the frequency tracking algorithms for random walk frequency fluctuations for (a) real- and (b) complex-valued signals.

of the band-pass filter. This can be a problem for a periodic component whose instantaneous frequency varies very rapidly. Indeed, whenever the estimate is not close enough to the true frequency, the underlying oscillation may be only extracted partially or even, in the worst case, not extracted at all. However, this corresponds to extreme cases as the convergence rate of the SFT is amply sufficient in most situations. Furthermore, when processing real data, the presented adaptive scheme is applied to the analytic representation of the input signal. Consequently, it is not applicable in real-time because of the block processing of the DHT. It is true that the DHT can also be approximated with a finite impulse response (FIR) filter, but this introduces a delay which is not always acceptable. A possible coarse solution to this problem is to apply the SFT directly to the real signal at the cost of slightly degraded performance. Another limitation of this technique is that it can only track one component at a time. Nevertheless, it should fail gracefully when applied to a signal with multiple oscillations by extracting one of the frequency components. In addition, it should be noted that one of the extensions for the SFT presented in Chapter 4 is designed to take into account such cases.

Extensions for the Single Frequency Tracker

4

In this chapter, three extensions to the single frequency tracker (SFT), presented in Chapter 3, are introduced. The motivations for extending the original adaptive algorithm are briefly summarized in Section 4.1. Then, the extensions for multiple frequencies, multiple signals and harmonic components are described in Sections 4.2, 4.3 and 4.4. Finally, their properties are discussed in Section 4.5. These extensions were presented in several publications [63, 64, 67–69].

4.1 Introduction

The SFT was presented in the previous chapter for an input signal composed of a complex cisoid with additive zero-mean noise. Although this setting is very simple, this adaptive algorithm is still able to extract an oscillation and its instantaneous frequency in more complex scenarios, such as in the presence of spurious periodic components. However, in many practical situations, the basic assumption of a single cisoid embedded in noise is too restrictive. Indeed, several important oscillatory components might be active simultaneously in the same signal. For instance, this is the case in electroencephalographic (EEG) data [70] where oscillations are present in different frequency bands. In such circumstances, the SFT is only able to track one of these components, and thus relevant information might be lost. Classic approaches, such as filterbanks and wavelet analysis, have already been applied in order to extract multiple oscillations at the same time. However, these methods implicitly assume that the instantaneous frequencies of the periodic components under study do not change over time. This can lead to the erroneous extraction of two successively-occurring components in adjoining bands where there is in fact only one oscillation with time-varying frequency. This problem raises the need for techniques capable of extracting multiple oscillatory components with time-varying instantaneous frequencies at the same time. One solution is to extend the SFT to the multiple frequency tracker (MFT) in order to take into account signals with several oscillations.

The single cisoid assumption of the SFT might also be overly optimistic in the closely-related problem of harmonic frequency components. For instance, harmonic components have been used to detect problems in induction motors [71, 72]. In such cases, at best, the estimation variance of the SFT will increase, but, at worst, the adaptive scheme will not track the fundamental frequency. A more clever approach for this type of signals would use the additional information provided by the harmonic components in order to achieve increased tracking performance and robustness. Indeed, the frequencies of harmonic components are integer multiples of the fundamental frequency. The harmonic frequency tracker (HFT), an extension to the SFT, is capable of tracking the instantaneous fundamental frequency as well as extracting the corresponding har-

monic components.

In several practical situations, the information of interest is observed with multiple sensors. In particular, this is the case for EEG and electrocardiographic (ECG) data. The resulting signals are often highly correlated. Therefore, instead of tracking the instantaneous frequency separately in each signal, jointly processing the available data will improve the tracking performance, the convergence speed and the robustness of the adaptive algorithm [73, 74]. The multivariate frequency tracker (MVFT) extends the SFT to multiple signals in order to take advantages of these considerations.

Similarly to the SFT, the three extensions are presented within the complex signal framework. Nonetheless, they are still applicable to the analytic representations of real signals thanks to the discrete Hilbert transform (DHT). And reverting back to real-valued data is as easy as taking the real part of the extracted oscillations. The additional bias and variance in the estimation caused by the DHT should remain negligible with sufficiently narrow band-pass filters. The MFT and MVFT are presented in Sections 4.2 and 4.3 respectively. Then, the HFT is described in Section 4.4. The MVFT is introduced before the HFT because the latter is based on the multivariate extension.

4.2 Multiple Frequency Tracker

The SFT is extended to track several periodic components simultaneously. Therefore, it is now assumed that the input signal is composed of K cisoids embedded in complex noise:

$$x[n] = \sum_{k=1}^K c_k[n] + v[n] = \sum_{k=1}^K A_k e^{j\omega_k n} + v[n], \quad (4.1)$$

where A_k and ω_k are the complex amplitudes and frequencies of the cisoids, and $v[n]$ is the additive complex noise. The basic idea behind the extension is to use one SFT to track each periodic component. However, because the band-pass filters (3.2) are not ideal ones (in fact, with an ideal filter no tracking would be possible), each of them will “see” the other components. This can lead to errors, especially for oscillations close in the frequency domain. For instance, several band-pass filters might collapse on the same oscillation. A solution to this problem is to use all-zero filters to cancel interference from other components as proposed by Rao and Kumaresan [16] in the context of speech processing. Therefore, the outputs of the MFT are the oscillations extracted by the cascade of all-zero filters and band-pass filters, $y_k[n]$ ($k = 1, 2, \dots, K$), and their estimated instantaneous frequencies. The structure of the extended tracking algorithm is summarized in Figure 4.1.

The number of oscillations K is assumed to be known. Each periodic component is tracked with two adaptive filters: an all-zero filter and a band-pass filter. The band-pass filter is the same as for the single frequency case (3.2), and thus the transfer function of the k th filter is given by

$$G_k(z; n) = \frac{1 - \beta}{1 - \beta e^{j\omega_k[n]} z^{-1}}, \quad (4.2)$$

where $\omega_k[n]$ is the instantaneous frequency estimate for the k th oscillation. The all-zero filter is composed of the $K - 1$ complex zeros whose locations are defined by the frequencies of the other components. However, since the true frequencies are unknown, they are replaced by the current estimates given by the SFTs. Consequently, the transfer function of the k th all-zero filter is

$$N_k(z; n) = C_k[n] \prod_{\substack{l=1 \\ l \neq k}}^K (1 - \gamma e^{j\omega_l[n]} z^{-1}), \quad (4.3)$$

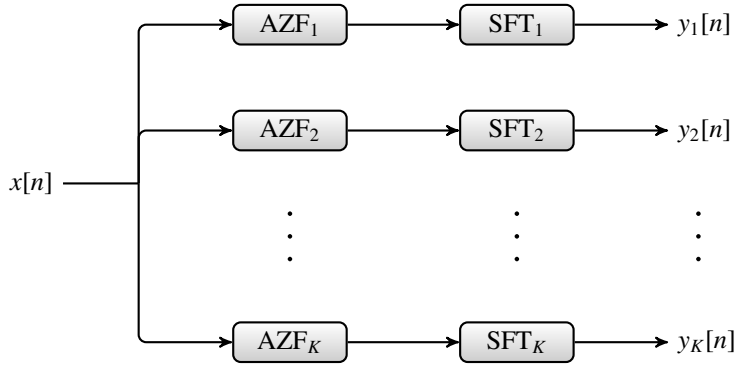


Figure 4.1: Structure of the MFT. Each periodic component is extracted with two stages, an all-zero filter (AZF) and a single frequency tracker (SFT).

where γ ($0 \ll \gamma \leq 1$) is the modulus of the zeros and $C_k[n]$ is a normalization coefficient to ensure unit gain and zero phase at frequency $\omega_k[n]$. The modulus parameter can be smaller than one, but it is typically set to one in order to suppress most of the other oscillations.

Although the additional all-zero filters cancel the interfering periodic components, they also introduce perturbations in the global response, which might not be purely band-pass anymore, resulting in unwanted frequencies leaking to the output. Moreover, longer filters reduce the convergence rate of the algorithm. An alternative is to select only one zero for the k th all-zero filter. This zero is placed at the tracked frequency closest to the k th frequency. This solution offers a good tradeoff between interference suppression, tracking speed and the band-pass characteristic of the overall filter. Furthermore, in situations where the tracking speed is not critical, another possibility is to use longer notch filters that result in better overall frequency response. For instance, this could be the case when extracting periodic components with slowly changing frequencies. Using a cascade of the notch filters described in Section 2.2.6 [56], the canceling filter for the k th component is defined as

$$N_k(z; n) = C_k[n] \prod_{\substack{l=1 \\ l \neq k}}^K \frac{1 - e^{j\omega_l[n]} z^{-1}}{1 - \lambda e^{j\omega_l[n]} z^{-1}}, \quad (4.4)$$

where λ ($0 \ll \lambda < 1$) is the pole modulus that control the notch width and $C_k[n]$ ensures unit gain and zero phase at $\omega_k[n]$ as before. An example of the global magnitude responses of the MFT is shown in Figure 4.2 for the three types of all-zero filters with the following parameters: $K = 3$, $\omega_1 = 0.2\pi$, $\omega_2 = 0.4\pi$, $\omega_3 = 0.8\pi$, $\beta = 0.95$, $\gamma = 1$ and $\lambda = 0.95$.

4.2.1 Multiple Frequency Tracking Example

The tracking capabilities of the MFT were tested on a synthetic signal of 2000 samples. This signal was composed of three cisoids with time-varying instantaneous frequencies embedded in additive complex circular Gaussian white noise. The first periodic component was a cisoid with a shift in frequency from 0.2π to 0.4π after 200 samples. The instantaneous frequencies of the last two cisoids changed linearly with time from 0.6π to 0.8π and from 0.9π to 0.1π . It is clear from the definition of the three components that their instantaneous frequencies intersected on several occasions. The noise level was set to 20 dB with respect to each cisoid. The MFT was applied with an all-zero filter for each interfering component. In other words, when tracking one of the cisoids the two others were canceled. Furthermore, the following parameters were selected for

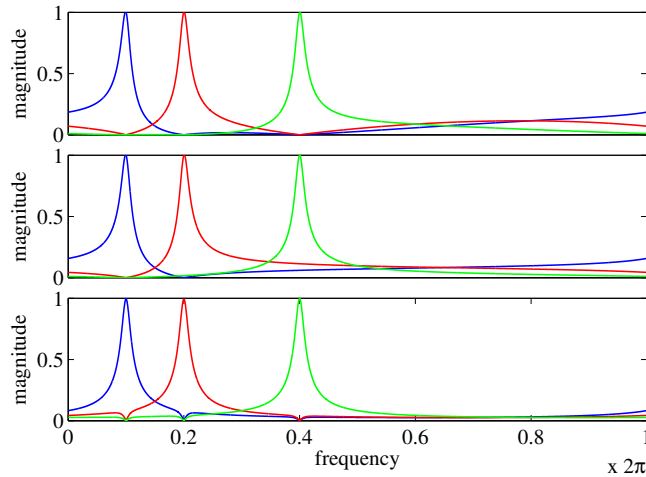


Figure 4.2: Global magnitude responses of the MFT when extracting three periodic components ($K = 3$) at $\omega_1 = 0.2\pi$ (blue), $\omega_2 = 0.4\pi$ (red) and $\omega_3 = 0.8\pi$ (green). The three types of canceling filters are shown, (top) each interfering component is suppressed with a zero (4.3), (middle) only the closest interfering frequency is canceled, (bottom) all the interfering oscillations are suppressed with a notch filter (4.4). The parameters of the filters are set as follows: $\beta = 0.95$, $\gamma = 1$, $\lambda = 0.95$.

the MFT: $\beta = 0.95$, $\delta = 0.95$ and $\gamma = 0.95$. And finally, the initial frequencies were set to 0.2π , 0.6π and 0.9π .

The instantaneous frequencies estimated by the MFT are shown in Figure 4.3 alongside the true values. Excluding the frequency shift where the algorithm needed some time to adapt, the frequencies of the cisoids were accurately estimated. Furthermore, this example illustrates that, with appropriate parameters, the MFT can even track components with time-varying frequencies that cross each other. The real parts of the input signal and the three complex oscillations extracted by the MFT are plotted in Figure 4.4. The cisoids were correctly filtered except when their instantaneous frequencies intersected and after the frequency shift.

4.3 Multivariate Frequency Tracker

The extension of the SFT to multiple signals is based on a generic approach [73, 74] applicable to a wide range of frequency tracking methods. Indeed, this approach was designed to easily extend adaptive algorithms for jointly tracking the instantaneous frequency of a periodic component present in several signals. Thus, the MVFT tracks the instantaneous frequency in L signals. These input signals are defined as

$$x_l[n] = c_l[n] + v_l[n] = A_l e^{j\omega_0 n} + v_l[n], \quad l = 1, \dots, L, \quad (4.5)$$

where A_l are the complex amplitudes, ω_0 is the frequency of the cisoids, and $v_l[n]$ are the mutually independent additive complex centered noises.

The approach for extending the SFT is to use the same band-pass filter (3.2) for the L signals. Then, the frequency of each cisoid, $\omega_l[n]$, is estimated in the same way as in the basic adaptive algorithm. Finally, the L frequency estimates are weighted in order to obtain a global estimate, $\omega[n]$. The structure of the MVFT is shown in Figure 4.5. The weights are chosen in order to

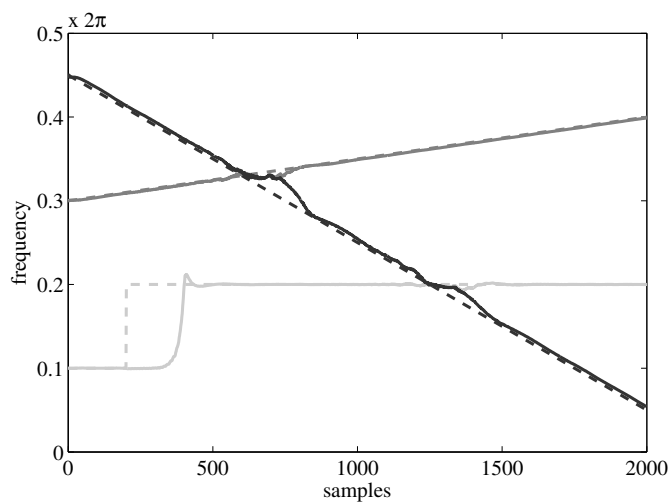


Figure 4.3: Instantaneous frequencies of three components estimated by the MFT (SNR = 20 dB, $\beta = 0.975$, $\delta = 0.95$, $\gamma = 0.95$). The estimated and true frequencies are denoted by solid and dashed lines respectively.

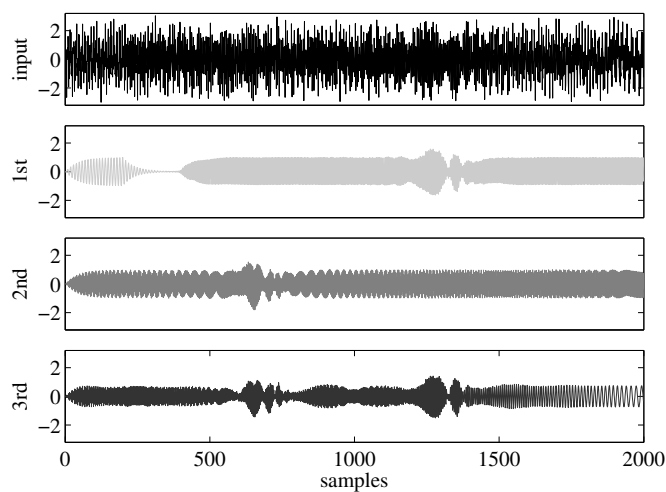


Figure 4.4: Real parts of the input signal and the three components extracted by the MFT (SNR = 20 dB, $\beta = 0.975$, $\delta = 0.95$, $\gamma = 0.95$).

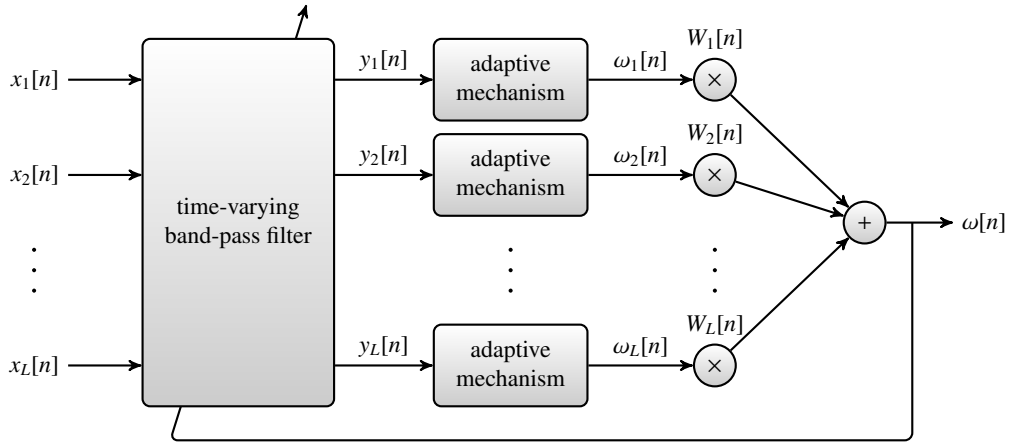


Figure 4.5: Structure of the MVFT. All input signals are filtered with the same time-varying band-pass filter. The instantaneous frequency of each extracted component is estimated independently and then the global estimate is obtained with a linear combination.

maximize a suitable measure of the tracking quality. For this purpose, the frequency estimates for the L signals at time n are rewritten as

$$\omega_l[n] = \omega_0 + \Delta_l[n], \quad l = 1, \dots, L, \quad (4.6)$$

where $\Delta_l[n]$ are the deviations from the true frequency caused by random perturbations. They are assumed to be zero-mean and uncorrelated. Therefore, the global frequency estimate is obtained as a linear combination of the estimates for all signals:

$$\omega[n] = \sum_{l=1}^L W_l[n] \omega_l[n] = \omega_0 + \sum_{l=1}^L W_l[n] \Delta_l[n], \quad (4.7)$$

with

$$0 \leq W_l[n] \leq 1, \quad l = 1, \dots, L, \quad \text{and} \quad \sum_{l=1}^L W_l[n] = 1.$$

Now, the goal is to find the set of weights that minimize the variance of this linear combination, or more precisely the variance of $\sum_{l=1}^L W_l[n] \Delta_l[n]$. This problem appears in financial portfolio optimization [75], and its exact solution, since the $\Delta_l[n]$ are assumed uncorrelated, is

$$W_l[n] = \frac{1/\sigma_{\Delta_l}^2[n]}{\sum_{i=1}^L 1/\sigma_{\Delta_i}^2[n]}, \quad l = 1, \dots, L, \quad (4.8)$$

where $\sigma_{\Delta_l}^2[n]$ is the variance of $\Delta_l[n]$. When $\sigma_{\Delta_1}^2[n] = \sigma_{\Delta_2}^2[n] = \dots = \sigma_{\Delta_L}^2[n]$, all the weights are equal to $1/L$. Otherwise, small variances lead to large weights. Thereby, the frequency estimates with low variance are favored.

The various steps of the MVFT are briefly summarized. First, all the input signals $x_l[n]$ are filtered with the same adaptive band-pass filter (3.2) in order to extract the L periodic components of interest $y_l[n]$:

$$y_l[n] = \beta e^{j\omega[n]} y_l[n-1] + (1-\beta)x_l[n], \quad l = 1, \dots, L. \quad (4.9)$$

Then, the instantaneous frequency of each extracted oscillation is estimated with the same mechanism as in the SFT (see Section 3.2):

$$\overline{Q}_l[n] = \delta \overline{Q}_l[n-1] + (1-\delta)y_l[n]\overline{y}_l[n-1], \quad l = 1, \dots, L, \quad (4.10)$$

$$\omega_l[n+1] = \arg\{\overline{Q}_l[n]\}, \quad l = 1, \dots, L, \quad (4.11)$$

where the overline denotes the complex conjugate.

From this point, a global estimate of the instantaneous frequency is computed by weighting the estimates obtained for each input signal according to (4.7). However, the variances of the deviations from the true frequency, $\sigma_{\Delta_l}^2[n]$, are not available in practice. And thus, the calculation (4.8) for obtaining the optimal values for the weights $W_l[n]$ is not applicable. Nonetheless, it is reasonable to suppose that the deviation variances are proportional to measures of the tracking quality, such as the cost function (3.6) of the algorithm which is based on the complex discrete oscillator equation. Consequently, it is assumed that, near convergence, the variances in (4.8) are proportional to the following ratios:

$$\sigma_{\Delta_l}^2 \propto \frac{J_l}{S_l}, \quad l = 1, \dots, L, \quad (4.12)$$

with

$$J_l = E\{|y_l[n] - e^{j\omega_l[n]}y_l[n-1]|^2\}, \quad l = 1, \dots, L, \quad (4.13)$$

and

$$S_l = E\{|y_l[n]|^2\}, \quad l = 1, \dots, L. \quad (4.14)$$

The output variances are used to yield a scale-independent scheme. In practice, similarly to the adaptive mechanism of the SFT, these quantities are estimated with exponentially weighted averages [65]:

$$J_l[n] = \mu J_l[n-1] + (1-\mu)|y_l[n] - e^{j\omega_l[n]}y_l[n-1]|^2, \quad l = 1, \dots, L, \quad (4.15)$$

$$S_l[n] = \mu S_l[n-1] + (1-\mu)|y_l[n]|^2, \quad l = 1, \dots, L, \quad (4.16)$$

where μ ($0 \ll \mu < 1$) is a forgetting factor. In most cases, μ is set equal to δ , the forgetting factor for frequency estimation. Monte Carlo simulations were used to confirm the validity of the proportionality assumption (4.12). The input signals were cisoids at frequency 0.4π with uniformly distributed random phase embedded in complex circular Gaussian white noise. The SNR values ranged from 0 to 20 dB in 1 dB steps. A modified SFT with a band-pass filter whose central frequency was fixed to 0.4π was then applied for frequency estimation. The deviation variance $\sigma_{\Delta_l}^2$ and the ratio J_l/S_l were computed over the last 1000 samples of 2000 iterations and averaged over 10,000 runs. The parameters of the adaptive scheme were set to $\beta = 0.95$ and $\delta = \mu = 0.95$. The ratio is plotted versus the deviation variance for all tested SNR levels in Figure 4.6 alongside a linear least squares fit. The Pearson coefficient of correlation [76] was significantly different from zero ($p < 10^{-5}$). Based on the results of the Monte Carlo simulations, it is clear that the proportionality assumption is indeed valid. Consequently, the weights used in practice are

$$W_l[n+1] = \frac{S_l[n]/J_l[n]}{\sum_{i=1}^L S_i[n]/J_i[n]}, \quad l = 1, \dots, L. \quad (4.17)$$

And the global estimate of the instantaneous frequency is

$$\omega[n+1] = \sum_{l=1}^L W_l[n+1]\omega_l[n+1]. \quad (4.18)$$

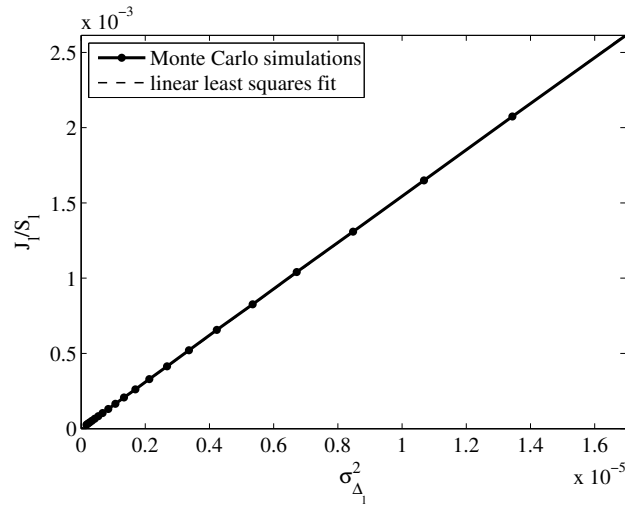


Figure 4.6: Validation of the proportionality assumption for the MVFT. The ratio J_l/S_l is plotted versus the deviation variance $\sigma_{\Delta_l}^2$ alongside a linear least squares fit. The cisoid frequency and the filter central frequency were set to 0.4π , and the algorithm parameters were set to $\beta = 0.95$ and $\delta = \mu = 0.95$.

4.3.1 Performance Analysis

A complete theoretical performance analysis of the MVFT is arduous due to the complexity introduced by the weighting procedure. In particular, closed-form expressions for the bias and variance of the estimated instantaneous frequency are very difficult to calculate. Nonetheless, in stationary conditions and near convergence, the instantaneous frequency estimate of the MVFT should be unbiased as it is a weighted sum of unbiased estimates (see Section 3.3.2). Moreover, when all input signals (4.5) have equal cisoid amplitudes and equal noise variances, the weights should be all equal to $1/L$. Consequently, the estimation variance of the MVFT should be decreased by a factor $1/L$ compared to the one of the SFT.

Monte Carlo simulations were used in order to validate these conjectures. Each input signal was composed of a cisoid at 0.4π with unit amplitude and uniformly distributed random phase embedded in complex circular Gaussian white noise. The SFT and MVFT were applied to estimate the instantaneous frequency for SNR values ranging from 0 to 20 dB in 1 dB steps. The multivariate extension was tested for two, three and four input signals. The same parameters were used for both tracking schemes, namely the bandwidth parameter was set to $\beta = 0.95$ and the forgetting factors were set to $\delta = \mu = 0.95$. The estimation bias and variance were computed over the last 1000 samples of 2000 iterations and averaged over 10,000 runs. Numerical simulations confirmed that the MVFT is indeed unbiased as shown in Figure 4.7. Also, the estimation variance of the MVFT is reduced with each additional input signal compared to the SFT as plotted in Figure 4.8. In fact, the gain factor in variance for the extended algorithm with respect to the basic scheme is approximately equal to the number of signals, as expected.

The convergence speeds of the SFT and MVFT were also compared with Monte Carlo simulations. As for the bias and variance, the input signals were cisoids with additive complex circular Gaussian white noise (SNR = 5 dB), with the exception that their instantaneous frequency changed abruptly from 0.2π to 0.4π . The SFT and MVFT for two and three signals were applied in order to estimate the instantaneous frequency. The bandwidth parameter was set to $\beta = 0.95$ for the three tracking schemes. However, different forgetting factors were used. Indeed, with the same parameters, the extended algorithm reduced the estimation variance. But, it is possible

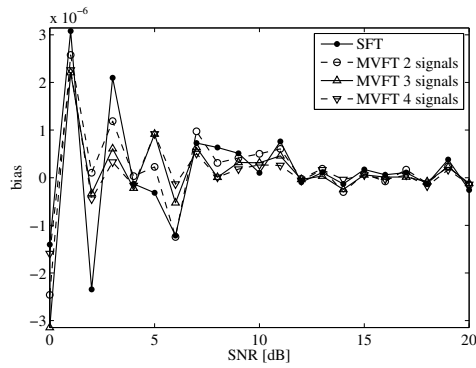


Figure 4.7: Bias of the SFT and MVFT for various SNR values ($A_l = 1$, $\omega_0 = 0.4\pi$, $\beta = 0.95$, $\delta = \mu = 0.95$).

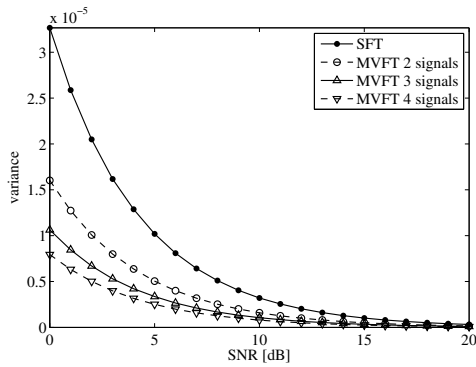


Figure 4.8: Variance of the SFT and MVFT for various SNR values ($A_l = 1$, $\omega_0 = 0.4\pi$, $\beta = 0.95$, $\delta = \mu = 0.95$).

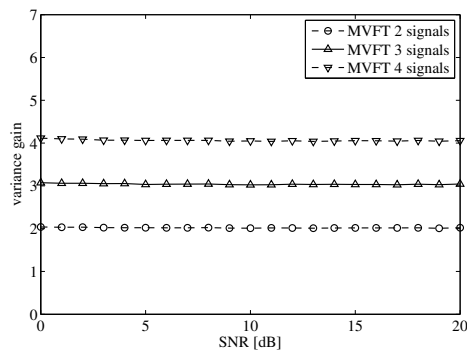


Figure 4.9: Variance gain of the MVFT with respect to the SFT for various SNR values ($A_l = 1$, $\omega_0 = 0.4\pi$, $\beta = 0.95$, $\delta = \mu = 0.95$).

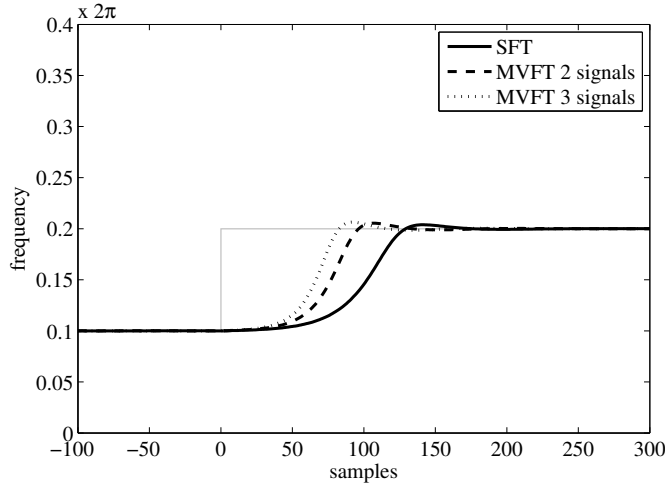


Figure 4.10: Convergence speed of the SFT and MVFT ($A_l = 1$, SNR = 5 dB, $\beta = 0.95$, $\delta = 0.95$ for SFT, $\delta = \mu = 0.9286$ for MVFT for two signals, $\delta = \mu = 0.9122$ for MVFT for three signals). The gray line denotes the true instantaneous frequency.

to choose the update parameters (or forgetting factors) so as to yield the same variance while increasing the convergence speed of the MVFT. To this end, these factors were set to $\delta = 0.95$ for the SFT, $\delta = \mu = 0.9286$ for the two-signal extension and $\delta = \mu = 0.9122$ for the three-signal extension. The values resulted in approximately the same estimation variance. The instantaneous frequency estimates averaged over 10,000 runs are shown in Figure 4.10 for the three adaptive algorithms. It is clear that using more signals resulted in increased convergence speed. Indeed, the MVFT for two signals was nearly 40 samples faster than the SFT, and the MVFT for three signals was about 60 samples faster.

4.3.2 Multivariate Multiple Frequency Tracker

The extensions to multiple frequencies and multiple signals can be straightforwardly combined in order to track several periodic components present in several signals. This combination is called the multivariate multiple frequency tracker (MVMFT). The input signals for this adaptive scheme are defined as

$$x_l[n] = \sum_{k=1}^K c_{l,k}[n] + v_l[n] = \sum_{k=1}^K A_{l,k} e^{j\omega_k n} + v_l[n], \quad l = 1, \dots, L, \quad (4.19)$$

where L is the number of signals, K is the known number of cisoids, $A_{l,k}$ and ω_k are the complex amplitudes and frequencies of the cisoids, and $v_l[n]$ are the mutually independent additive complex noises. The k th periodic component is extracted from the L input signals with the same cascade of an all-zero filter (4.3) and a band-pass filter (4.2) as in the MFT. This yields one output signal, $y_{l,k}[n]$, for each component and each input signal. Then, the SFT recursion is applied in order to compute an estimate of the instantaneous frequency of all extracted cisoids:

$$Q_{l,k}[n] = \delta Q_{l,k}[n-1] + (1-\delta)y_{l,k}[n]\bar{y}_{l,k}[n-1], \quad l = 1, \dots, L \text{ and } k = 1, \dots, K, \quad (4.20)$$

$$\omega_{l,k}[n+1] = \arg\{Q_{l,k}[n]\}, \quad l = 1, \dots, L \text{ and } k = 1, \dots, K. \quad (4.21)$$

Finally, the global instantaneous frequency estimate of each oscillatory component is obtained as follows,

$$\omega_k[n+1] = \sum_{l=1}^L W_{l,k}[n+1] \omega_{l,k}[n+1], \quad l = 1, \dots, L, \quad (4.22)$$

where the weights are given by

$$W_{l,k}[n+1] = \frac{S_{l,k}[n]/J_{l,k}[n]}{\sum_{i=1}^L S_{i,k}[n]/J_{i,k}[n]}, \quad l = 1, \dots, L \text{ and } k = 1, \dots, K. \quad (4.23)$$

As for the MVFT, the cost function and output variance estimates of all extracted components are computed with exponentially weighted averages,

$$\begin{aligned} J_{l,k}[n] &= \mu J_{l,k}[n-1] + (1-\mu) |y_{l,k}[n] - e^{j\omega_k[n]} y_{l,k}[n-1]|^2, & l = 1, \dots, L \text{ and } k = 1, \dots, K, \\ S_{l,k}[n] &= \mu S_{l,k}[n-1] + (1-\mu) |y_{l,k}[n]|^2, & l = 1, \dots, L \text{ and } k = 1, \dots, K. \end{aligned}$$

4.3.3 Multivariate Frequency Tracking Examples

Two main features of the MVFT and its extension to multiple frequencies, the MVMFT, are illustrated in this section with synthetic signals. First, their ability to deal with changes in the noise power is demonstrated. Second, they are shown to extract correctly periodic components common to several signals despite interfering oscillations.

In the first example, the MVFT was applied to two input signals of 1000 samples composed of a cisoid at 0.2π embedded in complex circular Gaussian white noise. In the first signal, the SNR was set to -5 dB for the first 500 samples and to 10 dB for the second half. It was the opposite for the second signal. The instantaneous frequency estimated by the MVFT was compared to the ones estimated by the SFT applied to each signal separately. The same parameters were selected for both tracking schemes: the bandwidth parameter was set to $\beta = 0.95$ and the forgetting factors were set to $\delta = \mu = 0.95$. The frequency estimates are shown in Figure 4.11 alongside the weights used by the MVFT. The frequency estimates obtained with the SFT are clearly inaccurate in the low SNR regions. By contrast, the frequency estimated by the MVFT is accurate for the whole duration of the signals. Looking at the adaptive weights reveals that during the first 500 samples the second signal was heavily favored because of its higher SNR. Following the SNR change, this situation was reversed. Indeed, due to the time-varying weights, the MVFT favors the signal with the best tracking performance at each time.

In the second example, two cisoids whose frequencies changed linearly from 0.7π to 0.5π and from 0.1π to 0.3π were common to the two input signals of 1000 samples. Interfering periodic components were added to each signal separately: a cisoid at 0.2π for the first signal and a cisoid at 0.6π for the second one. All cisoids (common and interfering ones) had the same amplitude. The MVMFT was applied in order to extract the two common oscillations simultaneously in both signals ($\beta = 0.95$, $\delta = \mu = 0.95$, $\gamma = 1$). For comparison purpose, the MFT was applied to each signal separately ($\beta = 0.95$, $\delta = 0.95$, $\gamma = 1$). Figure 4.12 shows the instantaneous frequencies estimated by both tracking schemes. The MVMFT correctly estimated the frequencies of the two common cisoids, despite the interfering components. However, it was not the case with the MFT. For the first signal, only the cisoid with frequency changing from 0.7π to 0.5π was tracked properly. The other cisoid was not extracted as soon as its frequency intersected with one of the interfering component. The second signal led to the opposite result. Therefore, the joint tracking of periodic components common to both input signals led to accurate frequency estimation. Furthermore, interfering components present in only one of the signals were discarded. This was not the case with separate tracking.

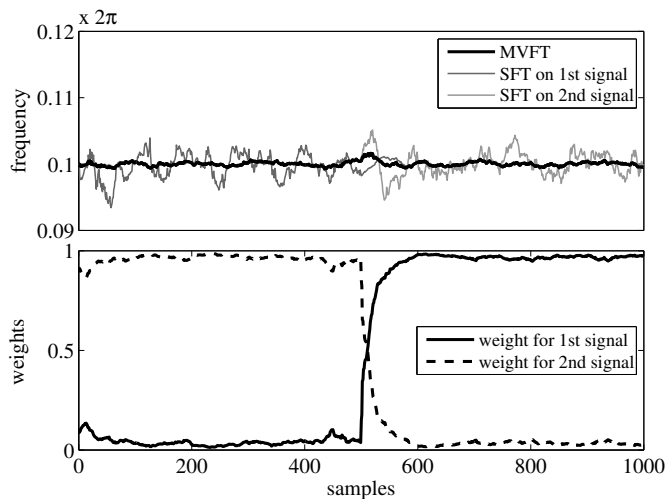


Figure 4.11: Instantaneous frequency and weights estimated by the MVFT ($\beta = 0.95$, $\delta = \mu = 0.95$). The estimates obtained when the SFT was applied to each signal separately are shown for comparison.

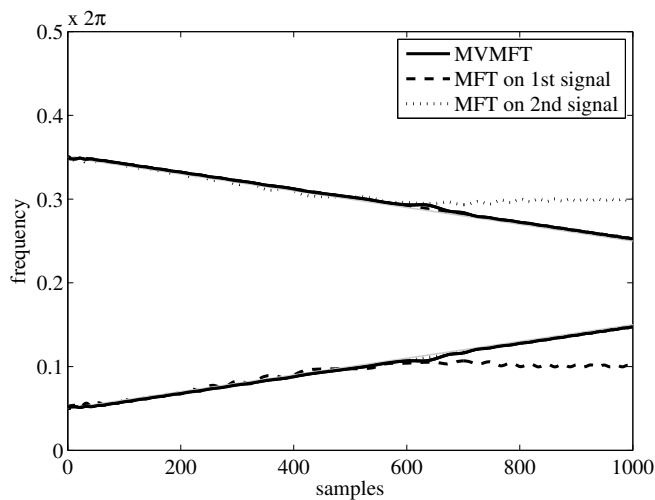


Figure 4.12: Instantaneous frequencies estimated by the MVMFT ($\beta = 0.95$, $\delta = \mu = 0.95$, $\gamma = 1$). The estimates obtained when the MFT was applied to each signal separately are shown for comparison. The gray lines denote the true instantaneous frequencies.

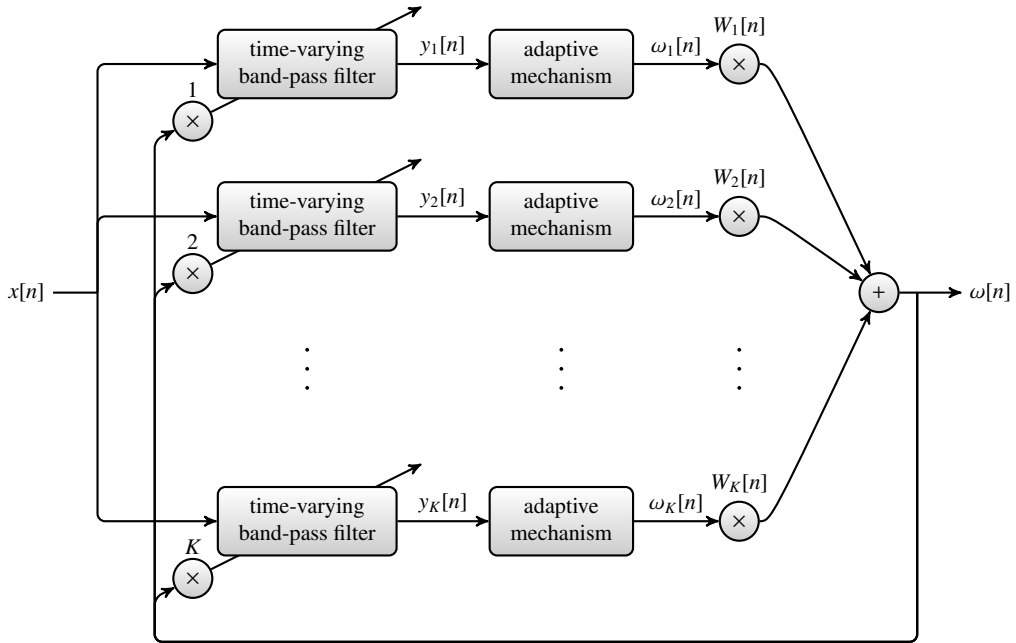


Figure 4.13: Structure of the HFT. The central frequencies of all time-varying band-pass filter are integer multiples of the fundamental frequency. The instantaneous frequency of each extracted component is estimated independently and then the global estimate is obtained with a linear combination.

4.4 Harmonic Frequency Tracker

The extension of the SFT to harmonic components, the HFT, is based on the two extensions presented previously. The idea of using one adaptive filter for extracting each harmonic component is derived from the MFT. Also, the same weighting procedure as in the MVFT is applied in order to compute an estimate of the fundamental frequency from the estimated harmonic frequencies. Therefore, each harmonic component is extracted with the same time-varying band-pass filter as in the SFT. Then, an estimate of the fundamental frequency is computed for each extracted oscillation with an adaptive mechanism. Finally, a linear combination is applied to obtain a global estimate. The structure of the HFT is depicted in Figure 4.13. By convention, the first harmonic component is considered to be the fundamental component throughout the description of this extended algorithm.

The input signal of the HFT is defined as

$$x[n] = \sum_{k=1}^K c_k[n] + v[n] = \sum_{k=1}^K A_k e^{jk\omega_0 n} + v[n], \quad (4.24)$$

where K is the known number of harmonic components, A_k are the complex amplitudes, ω_0 is the fundamental frequency, and $v[n]$ is the additive complex noise. The k th harmonic component is extracted with the same adaptive filter as in the SFT (3.2), except that the central frequency is k times the current estimate of the fundamental frequency, $\omega[n]$. The filtered outputs are given by

$$y_k[n] = \beta e^{jk\omega[n]} y_k[n-1] + (1-\beta)x[n], \quad k = 1, \dots, K. \quad (4.25)$$

From each of these extracted harmonic components, an instantaneous estimate of the fundamen-

tal frequency is computed using the oscillator-based recursion:

$$Q_k[n] = \delta Q_k[n-1] + (1-\delta)y_k[n]\bar{y}_k[n-1], \quad k = 1, \dots, K, \quad (4.26)$$

$$\omega_k[n+1] = \frac{\arg\{Q_k[n]\}}{k}, \quad k = 1, \dots, K. \quad (4.27)$$

As in the MVFT, a weighting procedure is applied to obtain a global estimate. For this purpose, instantaneous estimates of the cost function and output variance are computed with exponentially weighted averages for each extracted components:

$$J_k[n] = \delta J_k[n-1] + (1-\delta)|y_k[n] - e^{jk\omega[n]}y_k[n-1]|^2, \quad (4.28)$$

$$S_k[n] = \delta S_k[n-1] + (1-\delta)|y_k[n]|^2, \quad (4.29)$$

for $k = 1, \dots, K$. Using these two expressions, the weights are defined as follows,

$$W_k[n+1] = \frac{S_k[n]/J_k[n]}{\sum_{i=1}^K S_i[n]/J_i[n]}, \quad k = 1, \dots, K. \quad (4.30)$$

Finally, the estimate of the instantaneous fundamental frequency is

$$\omega[n+1] = \sum_{k=1}^K W_k[n+1]\omega_k[n+1]. \quad (4.31)$$

It should be noted that the $\omega_k[n+1]$ are estimates of the fundamental frequency since the harmonic frequencies are divided by k in (4.27).

The HFT has also two other interesting properties. First, it can track only a subset of the harmonic components. For instance, when the signal under study contains the first and third harmonic components, but no second component, the algorithm can be applied for only these components. Second, the weighting procedure can be easily extended for tracking the fundamental frequency in several signals with harmonic content as in the MVFT. Furthermore, when the harmonic frequencies are close, all-zero filters similar to the ones used in the MFT can be included in the HFT.

4.4.1 Performance Analysis

The same considerations regarding the complexity of a theoretical performance analysis of the HFT apply as for the MVFT. However, two remarks regarding the general behavior of this extension can be pointed out. First, the HFT is biased as the band-pass filters do not cancel the other harmonic components, resulting in leaked interfering frequencies. Nevertheless, the bias should remain negligible for sufficiently narrow filters. Second, the HFT should have reduced estimation variance compared to the SFT when tracking a signal with harmonic components as it uses more information.

These aspects of the extension to harmonic components were assessed by measuring the estimation bias and variance with Monte Carlo simulations. For this purpose, the input signals were defined as

$$x[n] = \sum_{k=1}^K \frac{1}{k} e^{j(k2\pi0.1n+\phi_k)} + v[n], \quad (4.32)$$

where the phase offsets ϕ_k were uniformly distributed between 0 and 2π , and $v[n]$ was a complex circular Gaussian white noise. The noise variance was computed with respect to the sum of the variances of all harmonic components. The SFT and HFT were compared for $K = 2$ and $K = 3$ with SNR values ranging from 0 to 20 dB in 1 dB steps. The adaptive parameters were

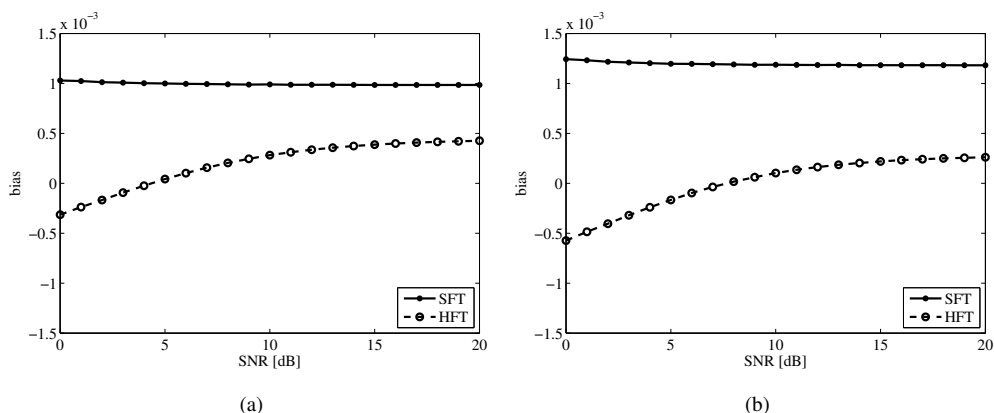


Figure 4.14: Bias of the SFT and HFT with respect to various SNR values ($A_k = 1/k$, $\omega_0 = 0.2\pi$, $\beta = 0.95$, $\delta = 0.95$), (a) for two harmonic components ($K = 2$), and (b) for three harmonic components ($K = 3$).

identical for the basic algorithm and its extension: $\beta = 0.95$ and $\delta = 0.95$. The estimation bias and variance were measured over the last 1000 samples of 2000 iterations and averaged over 10,000 runs. The results are shown in Figures 4.14 and 4.15. Both the SFT and HFT are biased when tracking a signal with harmonic components. This was expected for the basic adaptive scheme as the second and third harmonic components interfered with the estimation mechanism. The extended algorithm was also biased due to the fact that the other harmonic components were not canceled in the time-varying filters. Nonetheless, it should be noted that, in terms of absolute value, the bias of the harmonic extension is smaller for the two tested cases ($K = 2$ and $K = 3$). In any case, the bias of both methods can be considered negligible for the majority of applications. Concerning the estimation variance, the HFT yielded a more accurate estimate of the instantaneous fundamental frequency for all SNR values (with identical parameters). This decrease in variance was achieved using the additional information contained in all harmonic components.

Monte Carlo simulations were also run to compare the convergence speed of the SFT and HFT. The same input signal as before (4.32) was used with two harmonic components ($K = 2$), except that the fundamental frequency was abruptly shifted. Two cases were investigated: the fundamental frequency was changed from 0.2π to 0.3π and from 0.3π to 0.2π . The SNR was computed as before and was set to 5 dB. In addition, the bandwidth parameter was set to $\beta = 0.95$ for both the basic algorithm and its harmonic extension. For the same reasons as in the analysis of the convergence speed of the MVFT, different forgetting factors were selected in order to obtain approximately equal estimation variance for the SFT and HFT. Therefore, the SFT forgetting factor was set to $\delta = 0.95$ for the two cases, while it was set to $\delta = 0.9431$ when the fundamental frequency was shifted from 0.2π to 0.3π and to $\delta = 0.942$ when the frequency changed from 0.3π to 0.2π for the HFT. The estimates of the instantaneous fundamental frequency of the basic algorithm and its extension averaged over 10,000 runs are shown in Figure 4.16. When the frequency changed from 0.2π to 0.3π , the SFT converged faster than the HFT, albeit the latter has a smaller forgetting factor. Indeed, when the fundamental frequency was abruptly shifted, there were two periodic components with higher instantaneous frequencies that attracted the SFT. Whereas in the HFT, the adaptive mechanism responsible for estimating the frequency of the second harmonic component was closer to one of the first harmonic component directly after the shift. This slowed down the adaptation process of the harmonic extension compared to the

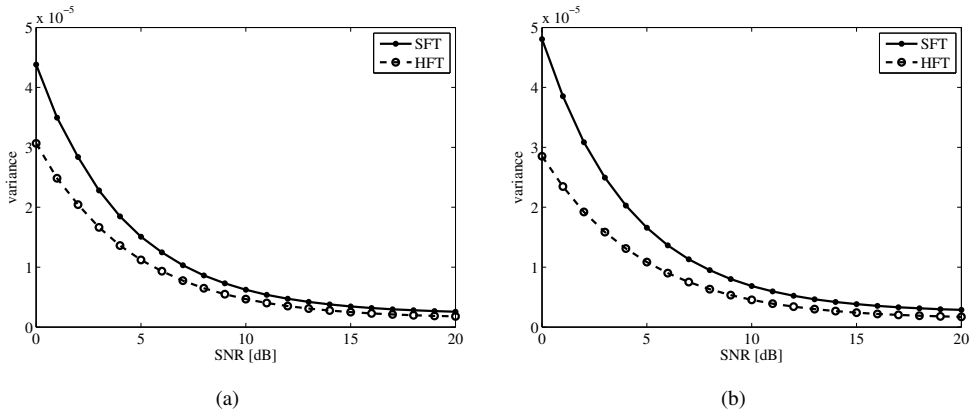


Figure 4.15: Variance of the SFT and HFT with respect to various SNR values ($A_k = 1/k$, $\omega_0 = 0.2\pi$, $\beta = 0.95$, $\delta = 0.95$), (a) for two harmonic components ($K = 2$), and (b) for three harmonic components ($K = 3$).

original method. By contrast, following the shift in fundamental frequency from 0.3π to 0.2π , the SFT faced the same problem and converged more slowly than the HFT. It is also worth mentioning that the number of samples required for convergence for the basic algorithm changed depending on the direction of the frequency shift (more samples were needed to converge after a decrease in frequency), while it remained roughly the same for the extension. In conclusion, despite being a little slower than the SFT in some situations, the HFT is attractive for signals with harmonic content as this extension can extract simultaneously every harmonic component of interest as well as estimate the instantaneous fundamental frequency with less pronounced bias than the original method.

4.4.2 Harmonic Frequency Tracking Example

The tracking capabilities of the HFT are illustrated with an example. The extended algorithm was applied to a signal composed of two frequency-modulated (FM) harmonic components embedded in noise. The input signal was defined as follows,

$$x[n] = \exp\{j2\pi 0.1n + j10 \cos(2\pi 0.002n)\} + 0.5 \exp\{j2\pi 0.2n + j20 \cos(2\pi 0.002n) + j\pi/4\} + v[n], \quad (4.33)$$

where $v[n]$ was a complex circular Gaussian white noise. The SNR value was set to 10 dB, and the extension parameters were set to $\beta = 0.95$ and $\delta = 0.95$. The frequency estimated by the HFT is shown on top of the true instantaneous fundamental frequency in Figure 4.17. One can observe that the estimate and the reference are in close agreement. Furthermore, as this tracking algorithm also extract the harmonic components, the latter are plotted in Figure 4.18. These components are readily available, without any additional filtering, when applying the HFT. And thus, they can be used directly in any further processing stages.

4.5 Discussion

In this chapter, three extensions for the SFT were presented and their advantages with respect to the original algorithm were discussed. Such extensions were developed because in many prac-

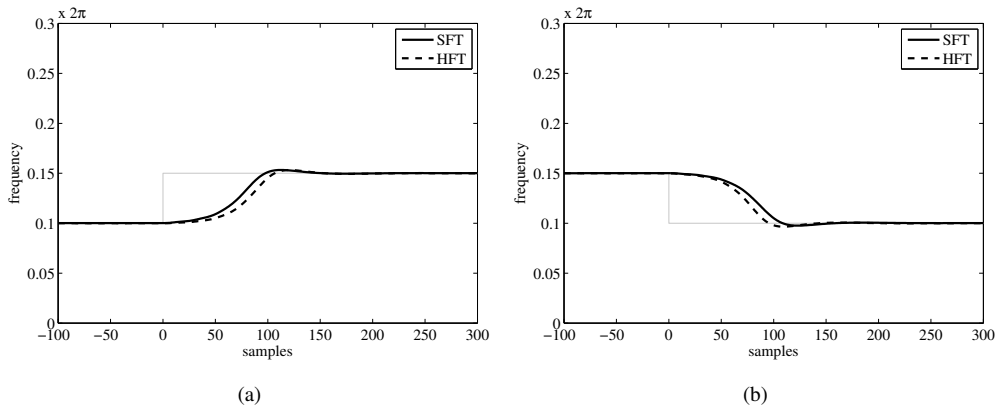


Figure 4.16: Convergence speed of the SFT and HFT ($A_k = 1/k$, SNR = 5 dB, $\beta = 0.95$, $\delta = 0.95$ for SFT), (a) frequency shift from 0.2π to 0.3π ($\delta = 0.9431$ for HFT), (b) frequency shift from 0.3π to 0.2π ($\delta = 0.942$ for HFT). The gray line denotes the true fundamental frequency.

tical situations the assumption of a signal composed of a single periodic component embedded in noise is far too restrictive. In particular, the case where a signal contains several simultaneous oscillations was investigated. A solution was proposed based on the basic idea of using one SFT to track each component. Filters were added in order to cancel interference between the adaptive mechanisms. This led to the MFT, an adaptive scheme capable of extracting separately periodic components and estimating their instantaneous frequencies. The performance of the SFT may also be insufficient when the signal under study contains harmonic components. This case is very similar to the previous one, except that there exists a structure in the oscillations. Indeed, the frequency of each harmonic component is an integer multiple of the fundamental frequency. Using this structure, the HFT was developed. This algorithm achieves lower estimation variance than the SFT while improving the robustness, as confirmed by Monte Carlo simulations. It also extracts the harmonic components which may be used directly in further processing.

Another case that was investigated is when the periodic component of interest is present in several signals. Indeed, nowadays many sensors are commonly used in an increasing number of applications. And while the SFT can be applied to one of the signals, this approach discards a lot of possibly relevant information. The original algorithm was thus extended in order to track the oscillation under study in several signals simultaneously. Monte Carlo simulations showed that the resulting adaptive scheme has better overall performance. Indeed, the MVFT achieved lower estimation variance and faster convergence speed, thanks to an adaptive weighting procedure. Moreover, this weighting procedure can also be included in the MFT and HFT so that they both benefit from the same advantages as the MVFT whenever more than one signal describing the phenomenon under study is available.

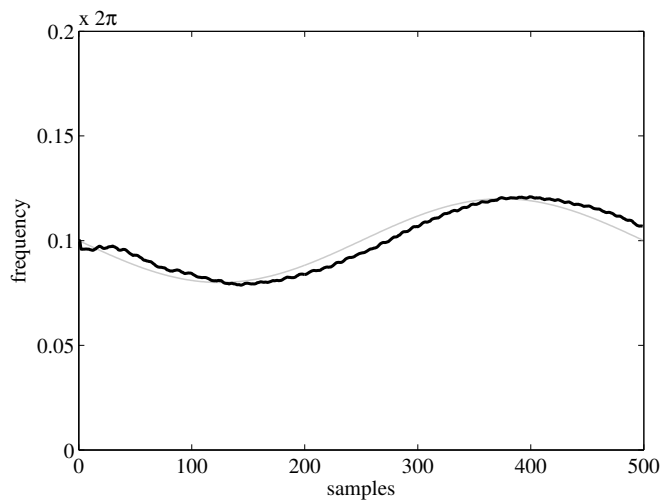


Figure 4.17: Instantaneous fundamental frequency of two FM harmonic components estimated by the HFT (SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$). The estimated and true frequency are denoted by black and gray lines respectively.

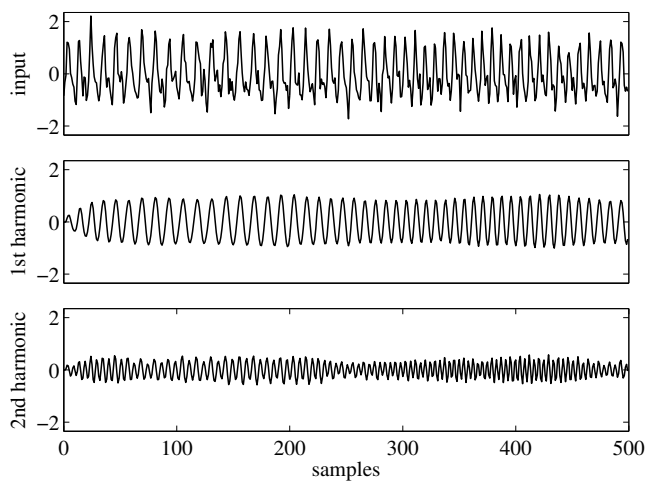


Figure 4.18: Real parts of the input signal and the two FM harmonic components extracted by the HFT (SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$).

Non-causal Instantaneous Frequency Enhancement

5

This chapter presents the non-causal frequency estimator (NFE), an algorithm for improving an initial frequency trajectory that can be obtained with any method. The grounds for developing the NFE are given in Section 5.1. Then, the algorithm is presented in detail in Section 5.2. Afterwards, the NFE performance is analyzed and compared to existing frequency estimation techniques in Section 5.3. A few examples of applications on synthetic and real data are given in Section 5.4. Finally, the merits and drawbacks of the presented method are discussed in Section 5.5.

5.1 Introduction

Frequency information is an important aspect in a wide range of fields such as biomedical engineering, speech processing and communications. In particular, instantaneous frequency estimation can shed light on a phenomenon evolving over time and can be used to extract precisely the underlying periodic component. Over the years, a large number of methods have been proposed for tracking the instantaneous frequency (e.g. [20, 35, 48, 51, 56, 77, 78]). The vast majority of these methods are causal, meaning that only past information is used for frequency estimation. While being a requirement for real-time applications, causality is not enforced for offline processing. Indeed, it seems intuitive that improved performance can be achieved using future information in addition to past information when estimating the instantaneous frequency. As a matter of fact, some methods, such as weighted finite differences of the phase [79], can be applied to sliding windows centered in time. Also, a few non-causal algorithms for frequency estimation have been proposed recently [80–83].

While using only past information, causal methods for frequency estimation can nonetheless achieve good tracking performance. But they introduce a delay in the estimation that can be quite severe after a sharp instantaneous frequency change. Therefore, given that causal methods yield rather satisfactory results in many situations, the NFE, a non-causal post-processing algorithm, was developed to update an initial estimate of the instantaneous frequency. Any method can be used for the initialization; it is not limited to a particular scheme. Furthermore, the proposed algorithm is specifically designed to estimate a frequency for which the oscillatory behavior in the input signal is maximal. Indeed, it uses a cost function derived from the complex discrete oscillator equation, like the single frequency tracker (SFT) [64] introduced in Chapter 3. As such, the NFE is presented within the complex-valued signal framework, but it is not restricted to complex signals. Using the discrete Hilbert transform (DHT) [34], it can be applied to the analytic representation of a real signal.

5.2 Algorithm

The input signal of the NFE is assumed to be a noisy cisoid,

$$x[n] = c[n] + v[n] = A_0 e^{j\omega_0 n} + v[n], \quad (5.1)$$

where A_0 and ω_0 are the complex amplitude and frequency of the cisoid, and $v[n]$ is an additive centered complex white noise. Moreover, it is supposed that N samples of the input signal are available ($n = 0, \dots, N-1$).

The algorithm is based on a cost function derived from the complex discrete oscillator equation. This equation is defined as

$$c[n] = e^{j\omega_0} c[n-1]. \quad (5.2)$$

It is verified for any cisoid at frequency ω_0 . Therefore, the frequency of a noiseless cisoid can be estimated with only two samples: $\omega_0 = \arg\{c[n]/c[n-1]\}$. When additive noise is present, frequency estimation must be performed more carefully. Therefore, the following local cost function was derived from the complex discrete oscillator equation,

$$J[n] = |x[n] - e^{j\omega[n]} x[n-1]|^2, \quad n = 1, \dots, N-1, \quad (5.3)$$

where $\omega[n]$ is the frequency estimate at time n . This estimate depends on time in order to take into account cisoids with time-varying instantaneous frequencies. Minimizing this cost function with respect to $\omega[n]$ yields the frequency trajectory for which the oscillatory behavior of the input signal is maximized.

The minimization of (5.3) is performed with a gradient descent algorithm. The update for the frequency estimate at each iteration and each time is defined as follows,

$$\omega_{i+1}[n] = \omega_i[n] - \mu \left. \frac{\partial J[n]}{\partial \omega[n]} \right|_{\omega[n]=\omega_i[n]} \quad n = 1, \dots, N-1, \quad (5.4)$$

where i is the iteration number, μ is the step-size, and $\omega_i[n]$ is the frequency estimate at time n and iteration i . This update is performed for all samples of the signal. The derivative of (5.3) with respect to $\omega[n]$ is given by

$$\frac{\partial J[n]}{\partial \omega[n]} = -2 \operatorname{Im} \{x[n] \bar{x}[n-1] e^{-j\omega[n]}\}, \quad n = 1, \dots, N-1, \quad (5.5)$$

where the upper bar denotes complex conjugation. It should be mentioned that complex numbers involved in this calculation do not lead to any difficulty since both $J[n]$ and $\omega[n]$ are real. Using a cost function derived from the discrete oscillator equation is advantageous compared to a classical approach in which the output power of a band-pass filter is used as the criterion for frequency estimation. Indeed, (5.5) yields a direction for the update whereas, in a power-based approach, both directions must be tested before updating the frequency estimate.

In practice, there are two main problems with the derivative in (5.4). First, the updates are very noisy since no averaging is performed. Secondly, there is no link between adjacent time instants, meaning that $\omega_i[n]$ could be completely different from $\omega_i[n+1]$. A simple yet effective solution is to smooth the derivative (5.5) over a window centered on time n . Therefore, a smoothed derivative is computed as follows,

$$d_i[n] = \sum_{k=-K}^K w[k] \cdot \operatorname{Im} \{x[n+k] \bar{x}[n+k-1] e^{-j\omega_i[n+k]}\}, \quad n = 1, \dots, N-1, \quad (5.6)$$

with $w[k]$ a smoothing window of length $2K+1$. The -2 factor in (5.5) is now included in the step-size μ . The window length must have odd length so it can be centered on a specific time

instant. Typically, a normalized Hamming window is a good choice for computing the smoothed derivative, but any meaningful window can be used. The samples at the borders ($n < K + 1$ and $n > N - K - 1$) are computed by mirroring the input signal. An important aspect of (5.6) is that the smoothing window is centered in time, meaning that both past and future data samples are considered. Therefore, the smoothing operation does not introduce any delay in the estimation of the derivative.

Finally, using (5.6), the update of the frequency estimate at each iteration is given by

$$\omega_{i+1}[n] = \omega_i[n] + \mu d_i[n], \quad n = 1, \dots, N - 1. \quad (5.7)$$

This iterative procedure is repeated until a stopping criterion s_i becomes smaller than a predefined threshold T . This criterion was selected to be the reduction in the normalized mean cost function after each iteration: $s_{i+1} = e_{i-1} - e_i$ with e_i defined as follows,

$$e_i = \frac{\sum_{n=1}^{N-1} |x[n] - e^{j\omega_i[n]} x[n-1]|^2}{\sum_{n=1}^{N-1} |x[n]|^2}. \quad (5.8)$$

This stopping criterion is scale-independent due to the normalization, and thus the threshold T does not depend the input signal power.

The input signal $x[n]$ is used directly for computing (5.6), but with a high level of noise, the update might become too noisy to be relied on. A solution is to band-pass filter the input signal in order to reduce the noise. However, the cisoid frequency can change over time. Therefore, a time-varying filter should be applied. More precisely, the central frequency of the band-pass filter should correspond to the true instantaneous frequency of the input signal at each time. But, since the true instantaneous frequency is unknown, it is replaced by its current estimate, $\omega_i[n]$. The band-pass filter is a Hamming window multiplied by a complex exponential. This filter was selected because of its simplicity, its finite impulse response, and because the Hamming window is the raised cosine window optimized to have a frequency response with minimal nearest side lobe [84]. The output of the filter at time n is thus given by

$$y_i[n] = \sum_{l=-L}^L h_{i,n}[l] \cdot x[n-l], \quad n = 0, \dots, N - 1, \quad (5.9)$$

where $2L + 1$ is the odd filter length and $h_{i,n}[l]$ is the time-varying band-pass filter. It depends on time and on iteration number as indicated by the subscripts. Furthermore, as for the computation of the smoothed derivative (5.6), the filtering operation does not introduce any delay in the output because it is centered in time and thus uses both past and future samples. The samples at the borders are computed by mirroring the input signal as for the smoothed derivative (5.6). The time-varying band-pass filter is defined as

$$h_{i,n}[l] = u[l] \cdot e^{j\omega_i[n]l}, \quad l = -L, \dots, L, \quad (5.10)$$

with $u[l]$ a normalized Hamming window of length $2L + 1$. The bandwidth of this filter depends directly on its length. It is also important to mention that it has zero phase and unit gain at the central frequency. Finally, once the filtered output $y_i[n]$ is computed, it is used instead of $x[n]$ in (5.6) and (5.8).

The implementation of the algorithm is straightforward, but a few aspects should be handled carefully. First, the power of the input signal could change over time. And the cost function and its derivative depend not only on the distance between the true frequency and its current estimate but also on the local power level. Therefore, estimation performance is influenced by the local power of the input signal. For example, a given step-size could yield a good update for the frequency estimate on some part of the signal while being too small or too large on another

Algorithm 5.1 Non-causal frequency estimator (NFE)

```

i = 0, s0 = ∞
set step-size  $\mu$ 
set threshold T for stopping the algorithm
initialize  $\omega_0[n]$  for  $n = 1, \dots, N - 1$ 
while  $s_i \geq T$  or  $s_i < 0$  do
  for  $n = 0$  to  $N - 1$  do
     $y_i[n] = \sum_l h_{i,n}[l] \cdot x[n - l]$ 
  end for
  for  $n = 1$  to  $N - 1$  do
     $p_i[n] = \sum_k w[k] \cdot |y_i[n + k]|^2$ 
     $d_i[n] = \sum_k w[k] \cdot \text{Im} \{ y_i[n + k] \bar{y}_i[n + k - 1] e^{-j\omega_i[n+k]} \}$ 
     $\omega_{i+1}[n] = \omega_i[n] + \mu(d_i[n]/p_i[n])$ 
  end for
   $e_i = (\sum_n |y_i[n] - e^{j\omega_i[n]} y_i[n - 1]|^2) / (\sum_n |y_i[n]|^2)$ 
   $s_{i+1} = e_{i-1} - e_i$ 
   $i = i + 1$ 
end while

```

part. Normalizing the smoothed derivative with a local estimate of the power leads to a scale-independent scheme. The local power estimate is computed with mirroring like the smoothed derivative (5.6),

$$p_i[n] = \sum_{k=-K}^K w[k] \cdot |y_i[n + k]|^2, \quad n = 1, \dots, N - 1. \quad (5.11)$$

Finally, using the local power estimate $p_i[n]$, the update for the frequency estimate becomes

$$\omega_{i+1}[n] = \omega_i[n] + \mu \frac{d_i[n]}{p_i[n]}, \quad n = 1, \dots, N - 1. \quad (5.12)$$

The length $2L+1$ of the time-varying band-pass filter (5.10) is an important parameter regarding the convergence of the algorithm since it controls the bandwidth. The noise will contaminate the frequency estimate if the filter is too wide. On the other hand, if it is too narrow, the algorithm will converge slowly or even not converge at all. An approach combining some advantages of both wide and narrow filters is to use a relatively wide filter (depending on signal properties) for estimating the instantaneous frequency and then use a narrow filter once the estimation has converged for precisely extracting the corresponding oscillatory component.

Another important aspect is the initialization of the algorithm. The initial frequency estimate can be obtained with any method. Even an estimate with non-negligible delay with respect to the true frequency can be used as a starting point. The step-size must be chosen carefully since it controls the performance of the algorithm. But due to the normalization by (5.11), it does not need to be adapted depending on the local power of the input signal. Usually, μ is chosen between 0.01 and 0.5. Similarly to the step-size, the threshold T for stopping the procedure is independent from the signal power because of the normalization. Typical values for T are between 10^{-9} and 10^{-6} , depending on the signal dynamics. The complete iterative procedure is summarized in Algorithm 5.1. It is also worth mentioning that the NFE can be straightforwardly extended in order to estimate the instantaneous frequency of a periodic component present in several signals by using a weighting scheme similar to the one presented in Section 4.3. Indeed, a weighted update approach [73] can easily be used to favor the signals with the most advantageous properties.

5.3 Performance Analysis

The performance of the NFE was analyzed with Monte Carlo simulations, because a theoretical analysis quickly becomes intractable. The NFE was compared to three existing non-causal techniques for frequency estimation: the short-time Fourier transform (STFT) [10], weighted finite differences of the phase (WFD) [79], and the adaptive notch smoother (ANS) [83]. Furthermore, the gain in estimation performance of the NFE with respect to a causal tracking scheme, namely the SFT presented in Chapter 3, was also investigated.

The input signal for the performance analysis of the four techniques was defined as

$$x[n] = e^{j\phi[n]} + v[n], \quad n = 0, \dots, N-1, \quad (5.13)$$

where $\phi[n]$ is the instantaneous phase and $v[n]$ is a complex circular Gaussian white noise. Two different cases were considered for the phase: constant frequency and linear chirps. Thus, the four methods were compared in stationary and non-stationary conditions. In the first case, the instantaneous phase was set to

$$\phi[n] = \omega_0 n + \theta, \quad n = 0, \dots, N-1, \quad (5.14)$$

where ω_0 is the frequency, and θ is an uniformly distributed phase term. In the second case, it was set to

$$\phi[n] = \begin{cases} \omega_1 n + \frac{\omega_2 - \omega_1}{N} n^2 + \theta & \text{for } n = 0, \dots, N/2 - 1, \\ (2\omega_2 - \omega_1)n + \frac{\omega_1 - \omega_2}{N} n^2 + (\omega_1 - \omega_2)\frac{N}{2} + \theta & \text{for } n = N/2, \dots, N-1, \end{cases} \quad (5.15)$$

where ω_1 and ω_2 are the frequency extrema, and θ is defined as before. These two instantaneous phases correspond to the following instantaneous frequencies:

$$\omega[n] = \omega_0, \quad n = 0, \dots, N-1, \quad (5.16)$$

and

$$\omega[n] = \begin{cases} \omega_1 + 2\frac{\omega_2 - \omega_1}{N}n & \text{for } n = 0, \dots, N/2 - 1, \\ 2\omega_2 - \omega_1 + 2\frac{\omega_1 - \omega_2}{N}n & \text{for } n = N/2, \dots, N-1. \end{cases} \quad (5.17)$$

The parameters for the Monte Carlo simulations were selected as follows: $N = 600$, $\omega_0 = 0.2\pi$, $\omega_1 = 0.2\pi$, $\omega_2 = 0.6\pi$. The instantaneous frequencies for these parameters are illustrated in Figure 5.1 for the two cases.

As the goal of this analysis was to compare as fairly as possible the STFT, WFD, ANS and NFE, the various parameters of these methods were chosen very carefully. The parameters of the NFE were selected first, and then appropriate values for the parameters of the other techniques were chosen. The smoothing window and filter lengths of the NFE were set to 91 ($K = 45$) and 21 ($L = 10$) respectively. The step-size was set to $\mu = 0.05$ and the threshold for the stopping criterion to $T = 10^{-7}$. In addition, a maximal number of iterations of 1000 was used, and thus the algorithm was stopped when the iteration counter i reached this value, even in the case it had not converged. Finally, the NFE was initialized to a constant instantaneous frequency: 0.2π in the first case and 0.4π in the second case. The STFT was computed on segments of 91 samples centered in time, and the shift between successive segments was set to one sample. Each segment was windowed with a Hamming window. All the discrete Fourier transforms (DFTs) were computed for 500 frequency points. The instantaneous frequency was estimated by searching, in each segment, the frequency bin with the largest magnitude. For the WFD method, the phase finite differences were calculated as

$$\Delta\phi[n] = \arg \{x[n]\bar{x}[n-1]\}, \quad n = 1, \dots, N-1. \quad (5.18)$$

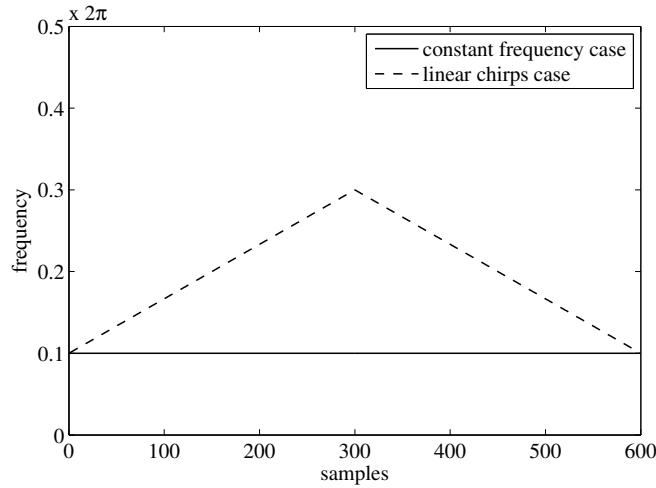


Figure 5.1: Instantaneous frequencies in the two considered cases for analyzing the NFE performance ($N = 600$, $\omega_0 = 0.2\pi$, $\omega_1 = 0.2\pi$, $\omega_2 = 0.6\pi$).

Then, these differences were divided into 91-samples overlapping segments shifted by one sample. Similarly to the STFT, these segments were centered in time. Finally, the instantaneous frequency estimate was obtained by smoothing the phase differences in each segment with Kay's weights [79]:

$$w[m] = \frac{3}{2} \frac{M}{M^2 - 1} \left(1 - \left(\frac{m - M/2 + 1}{M/2} \right)^2 \right), \quad m = 0, \dots, M - 1, \quad (5.19)$$

with $M = 91$. It is important to note that the same smoothing window/segment length was used for the NFE, STFT and WFD. The ANS is basically an adaptive notch filter that tracks an oscillatory component and its instantaneous frequency followed by a cascade of non-causal filters to compensate for the estimation delay. It requires three adaptation gains: $0 < \lambda < 1$ for amplitude rate adaptation which also determines the bandwidth of the filter used for extracting the periodic component, $0 < \gamma_\omega < 1$ for frequency adaptation, and $0 < \gamma_\alpha < 1$ for frequency rate adaptation, with $\gamma_\alpha \ll \gamma_\omega \ll \lambda$. According to [83], the gains for frequency and frequency rate adaptation can be defined on the basis of λ , and thus they were set to $\gamma_\omega = \lambda^2/2$ and $\gamma_\alpha = \lambda\gamma_\omega/4$. Two different values for the bandwidth parameter λ were chosen: $\lambda = 0.1$ which led to good estimation performance in high-noise scenarios, and $\lambda = 0.181672$ so that the band-pass filters used in the ANS and NFE had approximately equivalent 3-dB bandwidths. The initial frequency of the ANS was set to 0.2π in both cases.

The performance of each method was assessed by computing the mean squared error (MSE) between the estimated instantaneous frequency and the reference instantaneous frequency, (5.16) or (5.17). In order to avoid any border and initialization effects, the first and last 100 samples were discarded. The MSEs of the four techniques were averaged over 10,000 runs for SNR values ranging from -5 to 10 dB in 5 dB steps. They are shown in Figure 5.2 for the constant frequency case and Figure 5.3 for the linear chirps case. As there are large differences in scale across the different methods and SNR values, the MSEs are also reported in Table 5.1.

Looking at the outcomes of the Monte Carlo simulations, the first obvious result is that the WFD is the worst overall method for estimating the instantaneous frequency. Its performance quickly decreases for SNR values below 10 dB. Indeed, as it relies on the phase of the input signal for estimating the frequency, the WFD is very susceptible to broad-band noise [27]. For

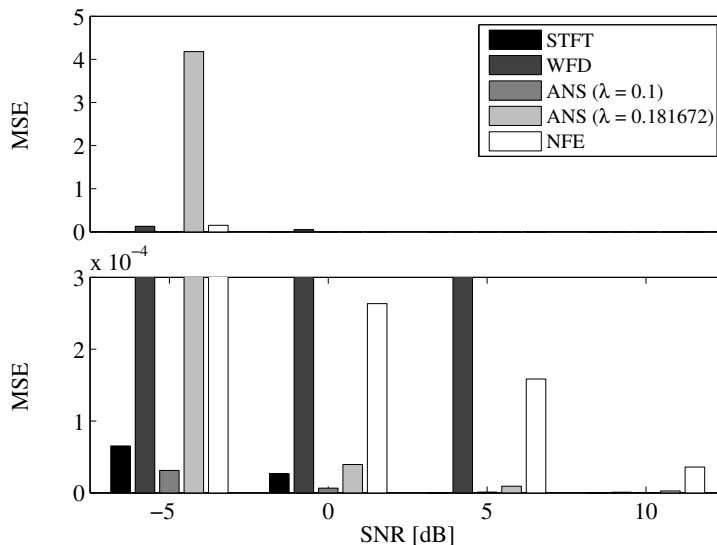


Figure 5.2: Mean squared errors of instantaneous frequency estimates for the constant frequency case. Top and bottom graphs show the same data at different scales.

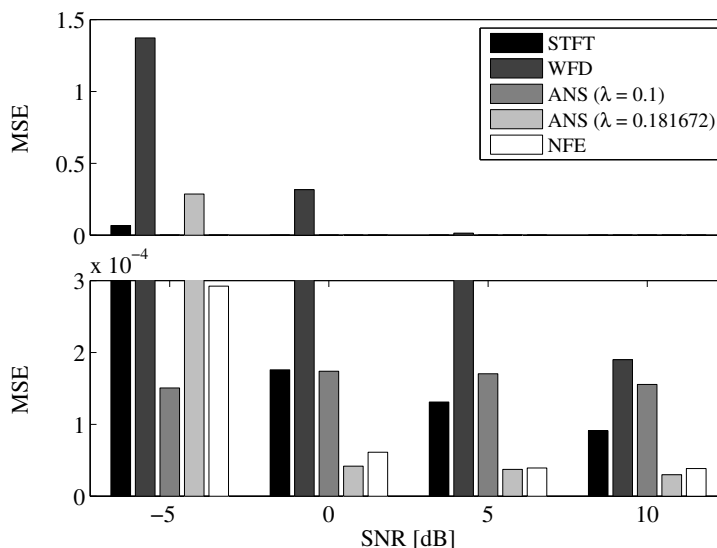


Figure 5.3: Mean squared errors of instantaneous frequency estimates for the linear chirps case. Top and bottom graphs show the same data at different scales.

Algorithm	-5 dB	0 dB	5 dB	10 dB
STFT	$6.52 \cdot 10^{-5}$	$2.66 \cdot 10^{-5}$	0	0
WFD	$1.30 \cdot 10^{-1}$	$5.19 \cdot 10^{-2}$	$2.35 \cdot 10^{-3}$	$9.00 \cdot 10^{-7}$
ANS ($\lambda = 0.1$)	$3.13 \cdot 10^{-5}$	$6.48 \cdot 10^{-6}$	$1.06 \cdot 10^{-6}$	$5.62 \cdot 10^{-7}$
ANS ($\lambda = 0.181672$)	4.18	$3.94 \cdot 10^{-5}$	$9.27 \cdot 10^{-6}$	$2.59 \cdot 10^{-6}$
NFE	$1.52 \cdot 10^{-1}$	$2.63 \cdot 10^{-4}$	$1.58 \cdot 10^{-4}$	$3.59 \cdot 10^{-5}$

(a) Constant frequency case.

Algorithm	-5 dB	0 dB	5 dB	10 dB
STFT	$6.73 \cdot 10^{-2}$	$1.76 \cdot 10^{-4}$	$1.31 \cdot 10^{-4}$	$9.11 \cdot 10^{-5}$
WFD	1.37	$3.16 \cdot 10^{-1}$	$1.42 \cdot 10^{-2}$	$1.90 \cdot 10^{-4}$
ANS ($\lambda = 0.1$)	$1.51 \cdot 10^{-4}$	$1.74 \cdot 10^{-4}$	$1.70 \cdot 10^{-4}$	$1.55 \cdot 10^{-4}$
ANS ($\lambda = 0.181672$)	$2.86 \cdot 10^{-1}$	$4.18 \cdot 10^{-5}$	$3.73 \cdot 10^{-5}$	$2.99 \cdot 10^{-5}$
NFE	$2.92 \cdot 10^{-4}$	$6.12 \cdot 10^{-5}$	$3.92 \cdot 10^{-5}$	$3.84 \cdot 10^{-5}$

(b) Linear chirps case.

Table 5.1: Mean squared errors of instantaneous frequency estimates.

the constant frequency case, the NFE remained behind the STFT and the ANS in terms of MSE performance, with the exception of the ANS with $\lambda = 0.181672$ for an SNR value of -5 dB which exhibited poor frequency estimation capabilities. The estimate provided by the STFT was particularly precise as the frequency of the input signal was set to an harmonic frequency of the DFT. However, for the linear chirps case, the NFE was the second best algorithm for all SNR values. And the ANS achieved the lowest MSEs for all noise levels only because two different values for the parameter λ were selected. Indeed, the ANS with $\lambda = 0.1$ was better than the NFE only when the SNR was set to -5 dB, while the latter led to lower MSEs for the other SNRs. Furthermore, although the ANS with $\lambda = 0.181672$ performed better than the NFE for SNR values greater than or equal to 0 dB, both methods were very close in terms of MSE. The STFT always lagged behind the NFE and ANS when the instantaneous frequency of the input signal varied linearly.

In addition to comparing the NFE to other non-causal techniques for frequency estimation, numerical simulations were also used to illustrate how this iterative method can achieve higher performance than causal algorithms such as the SFT. Clearly, this increased performance is not restricted to the NFE because, in general, the accuracy of frequency estimation can be improved by taking into account future information in addition to past information. The performance gain was assessed with a cisoid embedded in complex circular Gaussian white noise whose instantaneous frequency changed abruptly from 0.4π to 0.5π . Its phase followed a uniform distribution between 0 and 2π . The noise variance was set such that the SNR was equal to 5 dB. The various parameters of the SFT and NFE were carefully selected in order to compare both schemes as fairly as possible. Thus, the bandwidth parameter of the SFT and the half filter length of the NFE were set to $\beta = 0.818328$ and $L = 10$ respectively so that the time-varying filters of both algorithms had approximately equivalent 3-dB bandwidths. Moreover, the forgetting factor of the SFT was set to $\delta = 0.92$ while the half smoothing-window length and the step-size of the NFE were set to $K = 45$ and $\mu = 0.05$ respectively. This ensured that both techniques attained approximately equal levels of estimation variance after convergence. The initial frequency of the SFT was set to 0.4π . Whereas the NFE was initialized with the estimate given by the SFT. The instantaneous frequencies estimated by both methods were averaged over 10,000 runs. Figure 5.4 shows that the NFE provided a more accurate estimate of the frequency change than the

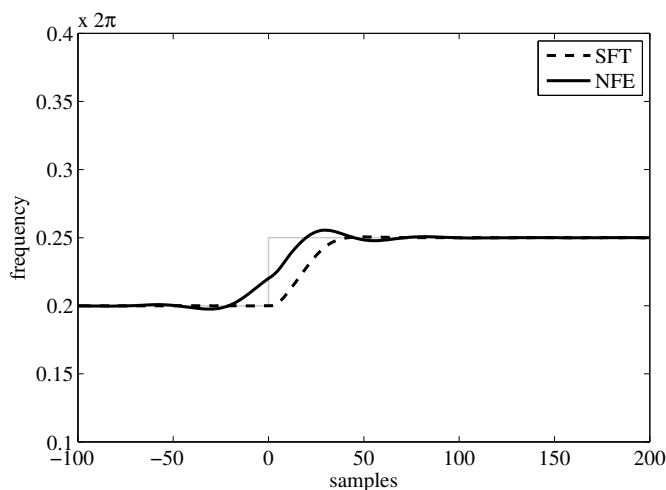


Figure 5.4: Convergence speed of the SFT and NFE. The instantaneous frequency estimates of both methods were averaged over 10,000 runs. The gray line denotes the true instantaneous frequency.

SFT. Clearly, this comparison is biased since the two methods did not use an equal amount of information. Nonetheless, it illustrates how frequency estimation performance can be improved by taking into account future information in addition to past information.

5.4 Examples

Four different signals were considered to illustrate the frequency estimation capabilities of the NFE. More precisely, two synthetic signals and two real signals were used. The first two highlighted how the NFE can handle oscillatory components with important changes in instantaneous frequency and enhance coarse initial frequency estimates. The third signal is an electroencephalographic (EEG) signal recorded during a visual evoked potential experiment, while the fourth one is an R-R interval signal. The R-R interval (or R wave to R wave interval) is the inverse of the heart rate, which is the number of heart beats per unit of time. Both the EEG and R-R interval signals exemplified the resilience of the NFE to interfering oscillations commonly observed in real data.

In the first example, the input signal was a cisoid whose frequency was modulated between 0.2π and 0.6π , with additive complex circular Gaussian white noise. The SNR was set to 0 dB and the frequency of the modulation was set to 0.004π . The initial frequency estimate was computed with the SFT ($\beta = 0.95$ and $\delta = 0.95$). The parameters selected for the NFE were the same as for the performance analysis ($K = 45$, $L = 10$ and $\mu = 0.05$). However, the cisoid was extracted with a narrower band-pass filter ($L = 40$) once the algorithm had converged. The instantaneous frequency estimates and the real part of the input and filtered output signals are shown in Figure 5.5. The frequency estimate given by the SFT is not really accurate. The algorithm is too slow to adapt to the fast-varying frequency. On the contrary, the estimate obtained after 286 iterations with the NFE is in close agreement with the true frequency. Also, the cisoid was precisely extracted with the time-varying band-pass filter. It is important to note that the SFT parameters were not optimized for this example. Indeed, it should be considered as an illustration of the estimation performance achieved by the NFE from an unsatisfactory starting point.

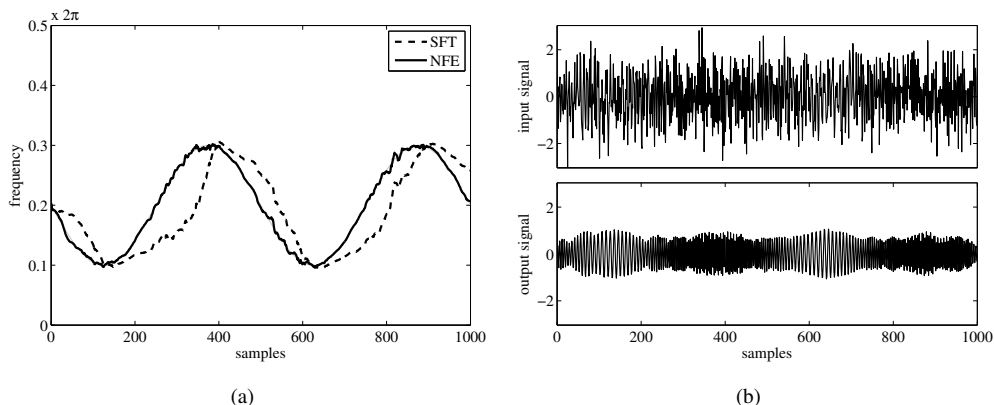


Figure 5.5: Frequency estimation for a frequency modulated cisoid embedded in white noise (SNR = 0 dB), (a) true frequency, SFT and NFE estimates, (b) real part of input and filtered output signals. The gray line denotes the true frequency.

The second example consisted of a linear chirp, whose frequency varied from 0.2π to 0.6π and back every 500 samples, embedded in additive complex circular Gaussian white noise. As in the previous example, the SNR was set to 0 dB. However, the NFE was initialized to a constant frequency of 0.4π for the whole signal duration. The purpose was to illustrate that the iterative algorithm can still converge when starting from a very coarse initial frequency estimate, such as the one obtained from a power spectral density estimate for instance. The same set of parameters as before was used. The instantaneous frequency estimate obtained after 310 iterations is plotted in Figure 5.6 alongside the real part of the input signal and the extracted periodic component. Clearly, the NFE correctly estimated the time-varying frequency despite the crude initialization to a constant frequency. As in the first example, the output signal was obtained with a narrower filter ($L = 40$) once the frequency estimate had converged.

The signal for the third example is an EEG signal recorded from electrode PO4 (parieto-occipital) during a visual evoked potential experiment with illusory contour stimuli [85]. This signal was chosen as it contained several interfering components in addition to the oscillation of interest. Before estimating the instantaneous frequency, the EEG signal was resampled at 250 Hz and band-pass filtered between 30 Hz and 60 Hz in order to discard the high-power content in the low frequencies. The initial frequency estimate was given by the SFT ($\beta = 0.975$ and $\delta = 0.95$), and the following parameters were used for the NFE: $K = 100$, $L = 30$ and $\mu = 0.05$. The parameters were set differently from the synthetic examples in order to follow as closely as possible the frequency changes and to take into account the simultaneous presence of several oscillations. The STFT and ANS presented in Section 5.3 were also applied in order to compare them to the NFE for processing a real EEG signal. However, the WFD was excluded as it exhibited the poorest performance with synthetic signals and due to its susceptibility to interfering oscillations. The STFT was computed on sliding windows of 101 samples centered in time so as to match the smoothing window of the NFE. The bandwidth parameter of the ANS was set to $\lambda = 0.1$ in order to estimate as precisely as possible the frequency of the main component. With higher values for this parameter, the ANS was attracted by spurious local oscillations and lost the main oscillatory component. The instantaneous frequency estimates computed by all these techniques are plotted on top of the STFT of the EEG signal in Figure 5.7. Clearly, the simultaneous presence of several oscillations showed the limits of the STFT, and the SFT needed some time to adapt to frequency changes as it used only past information. However, both the ANS

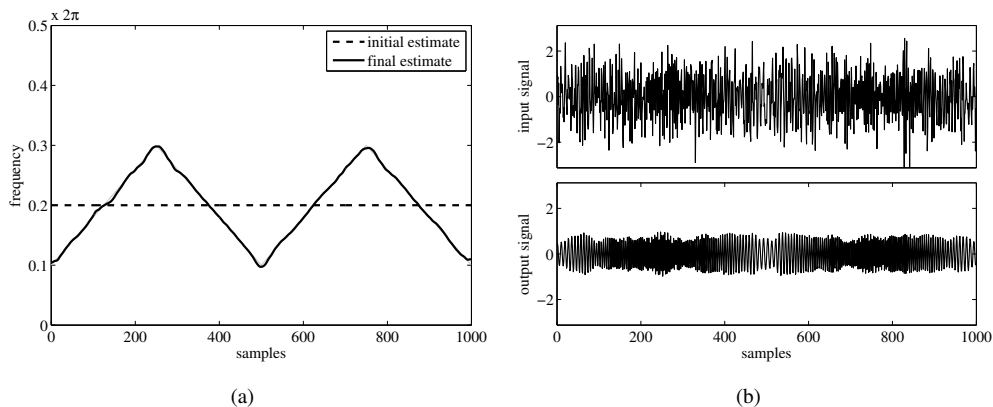


Figure 5.6: Frequency estimation for a complex linear chirp embedded in white noise (SNR = 0 dB), (a) true frequency, initial and final estimates, (b) real part of input and filtered output signals. The gray line denotes the true frequency.

and NFE yielded reliable estimates of the instantaneous frequency of the component between 30 Hz and 40 Hz. In particular, the ANS estimate was very smooth while the NFE estimate followed closely the slight frequency variations. It should be noted that, in this case, the ANS could not be made more reactive since as soon as the adaptation parameter λ was increased the estimate was disrupted by the spurious components with higher frequencies.

The R-R interval signal considered in the fourth example was recorded during the EuroBaVar study [86] aimed at comparing various techniques for measuring the baroreflex sensitivity. The R-R interval signals are typically used as measures of the heart rate variability, which is primarily caused by respiration and baroreflex in healthy subjects. These two causes are reflected as periodic components in R-R interval signals: a component around 0.1 Hz for the baroreflex and a component at the respiration frequency (around 0.25 Hz at rest) [87]. Since the signals from the EuroBaVar database were sampled beat by beat, the selected R-R interval signal was evenly resampled at 1 Hz. It was also high-pass filtered above 0.05 Hz in order to remove the very low frequency content with high power. Then, the SFT, ANS and NFE were applied in order to estimate the instantaneous frequency of the baroreflex. The STFT was not used in this case because it cannot handle signals with more than one oscillation. The parameters of the three techniques were selected in order to follow as closely as possible the component of interest while avoiding interference from the respiration oscillation. Therefore, the SFT was applied with $\beta = 0.95$ and $\delta = 0.975$, the ANS with $\lambda = 0.06$, and the NFE with the same parameters as for the EEG signal ($K = 50$, $L = 30$ and $\mu = 0.05$). The SFT and ANS were initialized to 0.1 Hz, while the initial estimate of the NFE was given by the SFT. The instantaneous frequency estimate obtained with these algorithms are plotted on top of the STFT of the considered R-R interval signal in Figure 5.8. Similarly to the previous example, the causal SFT yielded a delayed estimate that was not accurate constantly, particularly after sharp shifts in frequency. Again, the ANS provided a smooth estimate, but it failed to follow closely the main periodic component. Indeed, the selected value for λ was too small, but larger values resulted in the extraction of the respiration component around 0.3 Hz. By contrast, the NFE was able to estimate accurately the frequency variations of the baroreflex. Clearly, it can be argued that, for this example, the respiration component could have been filtered out in order to improve the estimation performance of all the techniques. But, in some practical situations, it is impossible to keep only the oscillation of interest. For instance, the frequency of the interfering component might be different for each signal of a database. This

would make difficult to choose an appropriate cut-off frequency.

5.5 Discussion

This chapter presented a non-causal algorithm for estimating the instantaneous frequency providing faster adaptation to sharp transitions than causal methods. This is made possible by using future information in addition to past information. While not suitable for real-time applications, the NFE can be used in every situation in which data processing is performed offline. Moreover, it is specifically designed to estimate the frequency of quasi-periodic components as its cost function is derived from the discrete oscillator equation. The NFE can also be initialized with an estimate obtained with any technique. One limitation however is that the bandwidth of the time-varying filter should ideally depend on the properties of the input signal. Also, since the NFE is an iterative algorithm, it is rather computationally intensive and might not always be adequate for processing large amounts of data. Nonetheless, it compared favorably with existing non-causal techniques for instantaneous frequency estimation on both synthetic and real data. In particular, the NFE proved to be resilient to interfering oscillations. And, as mentioned previously, it can be extended to estimate the instantaneous frequency of an oscillatory component present in several signals simultaneously in order to increase estimation robustness.

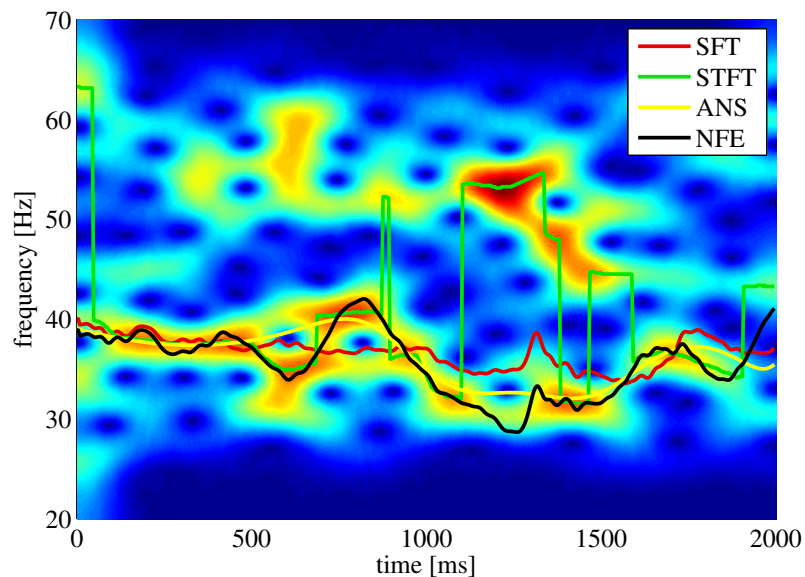


Figure 5.7: Estimation of the instantaneous frequency in an EEG signal with the SFT, STFT, ANS and NFE. The estimates are plotted on top of the STFT of the signal.

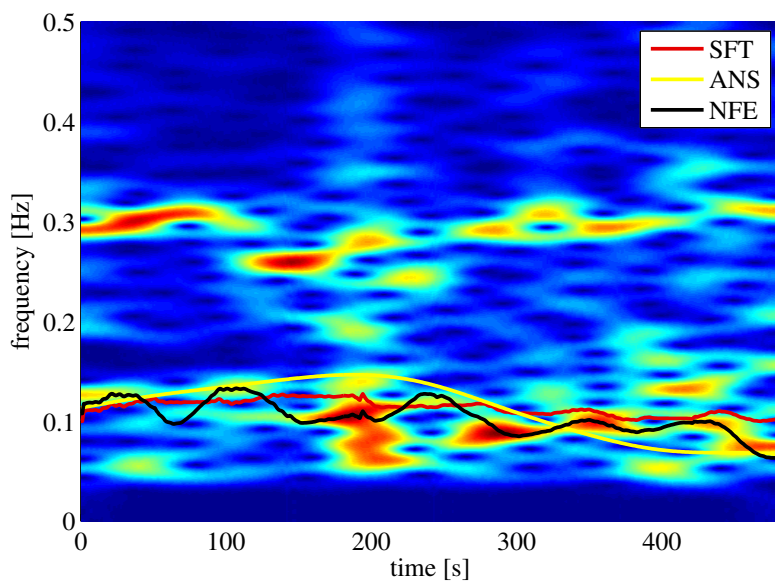


Figure 5.8: Estimation of the instantaneous frequency in an R-R interval signal with the SFT, ANS and NFE. The estimates are plotted on top of the STFT of the signal.

Part II

Frequency Tracking Applications

Adaptive Filters for the Identification of Couplings in Human EEG

6

An application of the single frequency tracker (SFT), introduced in Chapter 3, to electroencephalographic (EEG) data is presented in this chapter. The EEG is a recording of the brain electrical activity recorded from electrodes placed on the human scalp. It has been described as a “window on the brain” by Berger [88], the founding father of electroencephalography, since the features of EEG signals change depending on mental states. Furthermore, the scalp EEG provides large-scale measures of the brain activity, as a single electrode typically records the electrical action of tissue masses containing between 100 million and 1 billion neurons [89]. The EEG is used in both clinical and research contexts. In the former case, it is an important tool for following and treating certain illnesses, such as epilepsy, head trauma and coma [90]. In the latter case, it can help to highlight and understand cognitive processes such as motor process or short-term memory. An important aspect of EEG is its oscillatory behavior [70]. Indeed, neuronal oscillations occur in numerous situations, and traditionally they have been studied in different frequency bands [91]:

Delta band (0–4 Hz) Neuronal oscillations in this band have been observed in human sleep EEG [92]. They have also been found during continuous attention tasks [93].

Theta band (4–8 Hz) Activity in this band has been related to spatial processing [94]. It has also been shown to contribute to working [95] and long-term [96] memory processes.

Alpha band (8–12 Hz) Several roles have been proposed for alpha activity. Some studies suggested that it reflects top-down, inhibitory processes indirectly modulating cognitive processes [97], while others proposed that alpha oscillations have a direct function in neuronal processes such as memory, attention or awareness [98].

Beta band (12–30 Hz) Activity in this band has been observed in sensorimotor cortex in relation to motor behavior in humans [99]. It has also been shown that beta oscillations are involved in the maintenance of visual perceptions [100] and working memory [101].

Gamma band (30–70 Hz) Gamma oscillations have been observed in many brain regions, including the visual [102], auditory [103] and motor cortex [104]. A large body of research has implicated gamma activity in both bottom-up and top-down information processing (e.g. [105]).

Very high frequency band (>70 Hz) Although not part of the classical neuronal bands, many investigations suggest that very high frequency oscillations could have a significant role in cognitive processes. For instance, they have been reported in motor [106], somatosensory [107] and auditory [108] experiments.

These boundaries are very arbitrary and are defined slightly differently from one study to the other. Also, as most frequency bands are involved in a multitude of cognitive processes, only a few examples of such processes were given for each of these bands. Furthermore, more and more studies indicate that interactions or couplings between oscillations in different frequency bands play a key role in various cognitive processes [109].

This chapter illustrates how the SFT can help to identify these coupling mechanisms as well as to measure them more accurately. For this purpose, this adaptive frequency tracking algorithm was applied to scalp EEG data recorded during a visual evoked potential experiment. An introduction to brain oscillatory activity and couplings across different frequency bands is provided in Section 6.1. Then, the subject population, the acquisition protocol, the signal processing methods and the statistical analyses are described in Section 6.2. The results are presented in Section 6.3. Finally, the outcomes of this study are discussed in Section 6.4. An article [110] presents this study in detail.

6.1 Introduction

Oscillatory activity is a key component of brain dynamics and has been the focus of an increasing neuroscientific investigation. Neuronal oscillations have been considered a possible mechanism through which internal states exercise top-down influences on stimulus processing to impact perception [111, 112]. In particular, the phase synchronization of oscillatory components seems to be relevant for many cognitive processes [113]. Different models have been proposed for explaining the role of neural synchronization. For instance, the “communication through coherence” model [114] suggests that phase synchronization is a binding mechanism through which communication between different cortical areas is established. Another model proposes that phase synchronization facilitates neuronal plasticity [115]. Other studies [116, 117] consider that large-scale integration of perception into a unified representation is supported by neural synchronization. Therefore, synchronization of neuronal oscillations is considered a key mechanism for solving the problem of binding multiple and/or distributed representations. Moreover, this mechanism not only encompasses interactions between different cortical areas but also interactions between classical neuronal frequency bands; so-called cross-frequency couplings [109]. These cross-frequency couplings have been proposed as a framework for unifying the neuronal oscillations at different temporal and spatial scales [118]. The importance of these coupling processes have been demonstrated in recent studies of motor, sensory and cognitive tasks (e.g. [119–126]).

The reliability of methods for identifying these interactions across frequency bands can be examined using the well-known illusory contour (IC) stimuli [127]. Investigators have considered this paradigm as exemplary of the binding problem because physically absent borders of an object must be “filled-in” (at least perceptually if not also neurophysiologically) between inducers. One consistent observation is increased gamma power for IC vs. control stimuli (e.g. [128–130]). Another highly replicable finding is stronger global field power in the event-related potential (ERP) to the presence vs. absence of ICs (e.g. [85, 131–134]). The case of IC processing thus exemplifies a situation where the relationship between effects observed using analyses of event-related potentials (ERPs; which are heavily influenced by lower-frequency oscillations below ~ 25 Hz) and those obtained using time-frequency analyses (which typically focus on higher-frequency oscillations above ~ 25 Hz) remains to be detailed and ultimately conjoined (e.g. [135]). Moreover and despite being the subject of neuroscientific investigation spanning several decades in both humans and animal models, controversy persists regarding whether ICs are the result of bottom-up vs. top-down mechanisms (e.g. [134]). These kinds of results highlight the need for signal processing methods that can detail relationships between extracted features in a statistically sound manner.

Neural synchronization underlying cross-frequency couplings has been studied with a large number of different tools [136]. In particular, methods based on phase information, such as phase locking value [137, 138], have been applied to EEG data. Moreover, it has been shown recently that phase can encode more information than power [139], and thus such methods are well-suited to analyze cross-frequency interactions. The phase information is typically extracted with the widely-used Hilbert transform [23], but it should be considered with caution. The extracted phase is guaranteed to be physically meaningful only for narrow-band signals [27], and thus phase interpretation is problematic for broad-band signals. It should be noted that this interpretation problem arises with any technique for phase extraction. Consequently, the phase locking value is sensitive to broad-band interference [140]. A straightforward solution to this problem consists of adding a pre-processing step that separates EEG data into various narrow frequency bands with band-pass filters or wavelet analysis. Although this filter-bank approach can lead to more reliable analyses of cross-frequency couplings [119], it has a major disadvantage. The specifications of the filters (e.g. cut-off frequencies, attenuation, etc.) are predefined without taking into account the dynamics of the EEG signal under investigation. Therefore, an oscillatory component whose instantaneous frequency crosses the limit between two bands would be considered as two different oscillations occurring successively. In such cases it would be preferable to apply adaptive methods that can track a periodic component with a time-varying instantaneous frequency in a continuous manner. The SFT [63, 64], presented in Chapter 3, is such a technique in which a time-varying band-pass filter is adapted over time in order to extract an oscillation and its instantaneous frequency.

In this study, such adaptive filters are used to analyze the evolution of phase-amplitude and phase-phase couplings in response to the presence vs. absence of ICs. Although the separation of the signals into various frequency bands still relied on predefined band-pass filters, wider filters than the ones typically used in practice for processing EEG data were chosen. The following step was to retrieve the main oscillatory component and its instantaneous frequency in each band with the SFT. Thus, narrow-band signals were obtained from which the phase information could be precisely extracted, which, in turn, was used for measuring phase-amplitude and phase-phase coupling strength over time. The complete procedure is summarized in Figure 6.1. In more detail, three aspects of cross-frequency couplings during IC perception were tested. First, it was checked that stimuli with and without IC elicited a change in terms of coupling strength by using surrogate stationary signals generated from the original EEG data. This analysis assessed if the two types of stimuli caused a response before conducting further tests. Second, the responses to stimuli with IC were contrasted to the ones without such contours. One of the goal of this study was to determine if ICs elicited specific changes in terms of coupling strength. Last, the results obtained with and without adaptive frequency tracking were compared in order to highlight the value of the SFT for precisely extracting the phase information and measuring cross-frequency couplings. Synthetic signals were also used to show the advantages of frequency tracking. Overall, this study constituted a thorough analysis at a group-level as well as dynamic analysis of instantaneous frequency and cross-frequency couplings following IC presentation.

6.2 Materials and Methods

6.2.1 Experimental Setup

This study is a new analysis of a subset of data appearing in a previously published study that focused on broad-band ERPs in response to IC stimuli [85]. Full details regarding the experimental setup can be obtained from the original study. Here, only the essentials are provided. The participants included nine healthy adults (seven men and two women), aged 22–47 years (mean \pm standard deviation = 34 ± 10 years). Seven of the participants were right-handed and

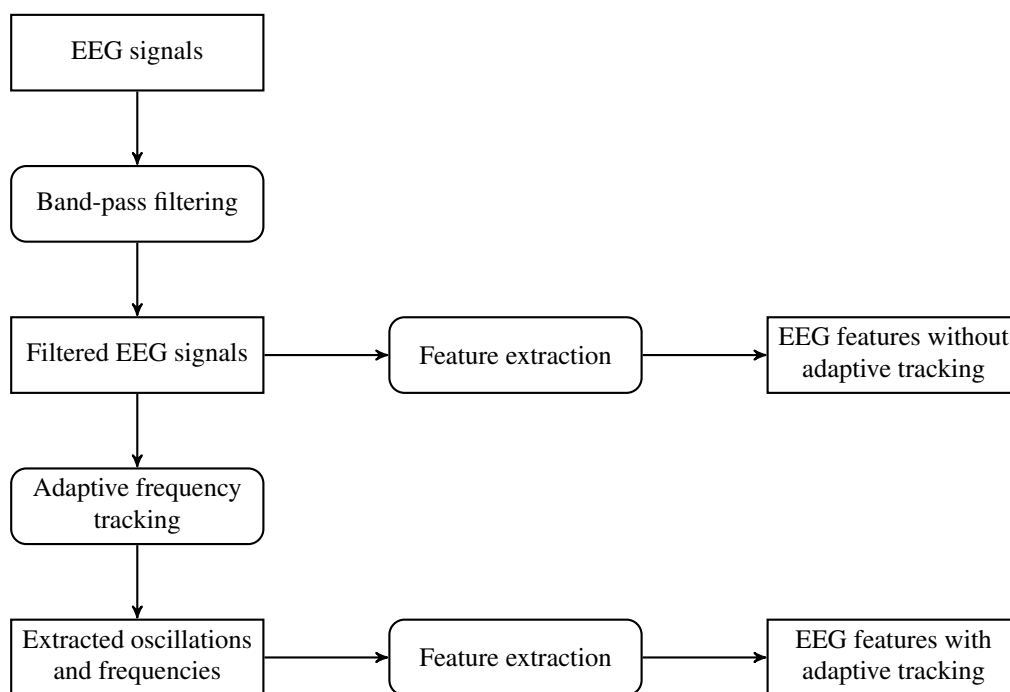


Figure 6.1: Processing steps used for analyzing the EEG data from the IC experiment. The same features were extracted before and after frequency tracking in order to assess the usefulness of the SFT.



Figure 6.2: Experimental conditions, (left) IC condition, (right) NC condition.

two left-handed according to the Edinburgh handedness inventory [141].

Participants viewed arrays of “pac-men” inducers presented in either of two orientations. In the illusory contour (IC) condition, the inducers were turned in order to produce the illusory perception of a simple geometric shape. On the contrary, in the no contour (NC) condition, the inducers were rotated 180° outwards; this prevented any illusory perception with the same luminance and contrast. Examples of the two conditions are shown in Figure 6.2. Each stimulus appeared for 500 ms, followed for 1000 ms by a blank screen. A Yes/No response prompt appeared then and remained visible until a decision was made, allowing subjects to control stimulus presentation. The response was followed by a blank screen for 1000 ms. The response prompt was used to clearly separate the sensory response from the motor response.

6.2.2 Data Acquisition and Pre-processing

Continuous EEG was recorded through a Neuroscan Synamp from 64 scalp electrodes (impedances $\leq 5 \text{ k}\Omega$), referenced to the nose, band-pass filtered from 0.05 Hz to 100 Hz, and sampled

at 500 Hz. Trials for each subject and condition were visually selected with the Cartool software by Denis Brunet [142] (<http://brainmapping.unige.ch/cartool>). A threshold of $\pm 80 \mu\text{V}$ for artifact rejection was used. Each trial represented 2000 ms of EEG data, with stimulus onset after 500 ms. There were an average (\pm standard deviation) of 300 (± 57) EEG trials from the IC condition and 295 (± 43) trials from the NC condition included in the analyses. These values did not significantly differ ($t_{(8)} = 1.06$, $p = 0.32$). Once the trials were extracted, all further processing was performed in MATLAB[®].

First, the EEG signals recorded from a cluster of five electrodes (P2, P4, P6, PO4 and PO6) were selected. This selection was based on the right-lateralized posterior scalp distribution of the ERP difference between IC and NC conditions (cf. Figure 3 in [85]). Then, signals from these electrodes were resampled from 500 Hz to 250 Hz and the power line interference at 60 Hz was canceled with a narrow notch filter. (The original recordings were performed at the Nathan Kline Institute for Psychiatric Research in Orangeburg, New York, USA.) The spatial mean of the five electrodes was computed in order to obtain a slightly more global view. The signals obtained after spatial averaging were then filtered in the following frequency bands: 1–4 Hz, 4–8 Hz, 8–12 Hz, 15–25 Hz, 35–45 Hz, 45–55 Hz, 55–65 Hz and 65–75 Hz. It is important to mention that the signals were filtered in both forward and reverse directions in order to achieve zero phase distortion [31]. Consequently, the samples at the beginning and at the end should be considered with caution as they were susceptible to transients and border effects. Once the band-pass signals were obtained, the SFT [63, 64] was applied in order to extract the main oscillation in each band as well as its estimated instantaneous frequency.

As the SFT is thoroughly presented in Chapter 3, only its main features are repeated here. It has been developed in the complex-valued signal framework, and thus it must be applied to the analytic representation of the input signal. This representation can be obtained with the discrete Hilbert transform (DHT) [34]. A real output signal can always be recovered by keeping only the real part. It is composed of two parts (Figure 3.1): a time-varying band-pass filter for extracting the oscillation and an adaptive mechanism for updating the current instantaneous frequency estimate. The filter (3.2) has unit gain and zero phase at its central frequency, which is of the utmost importance for measuring couplings based on phase information. The adaptive mechanism is based on the complex discrete oscillator equation (3.4), and thus the SFT maximizes the oscillator behavior at the output.

To summarize the pre-processing, the spatial average of a cluster of five electrodes was computed, then the signals were separated into various frequency bands with fixed band-pass filters, and finally the SFT was applied in order to extract the main oscillation and its instantaneous frequency in each band with a narrow time-varying band-pass filter. The following values were selected for the parameters of the tracking scheme: $\beta = 0.975$ (this corresponds approximately to a 3-dB bandwidth of 2 Hz), $\delta = 0.95$, and the initial frequency was set to the center of the considered frequency band. For proper initialization of the internal variable, the SFT was applied to longer signals obtained by adding the mirrored first 500 ms at the beginning. The input and outputs of the SFT are illustrated in Figure 6.3 with an EEG signal filtered in the 35–45 Hz band. It is possible to see that the tracking is not immediate, because the adaptation introduces a slight delay. Finally, one oscillatory component and its estimated instantaneous frequency were obtained for each frequency band.

In order to assess the presence or absence of phenomena, surrogate EEG signals were generated from the original data [143, 144]. This was done as follows: (1) an EEG signal was transformed into the frequency domain with the discrete Fourier transform (DFT), (2) then the amplitudes were kept but the phases were randomized (random variables drawn from a uniform distribution between 0 and 2π), (3) finally the modified signal was transformed back into the time domain. The phase randomization destroys the structure in the input signal and yields a more stationary output. However, the surrogate signal shares some properties with the original one, such as probability density function and autocorrelation. Thus, they have also the same

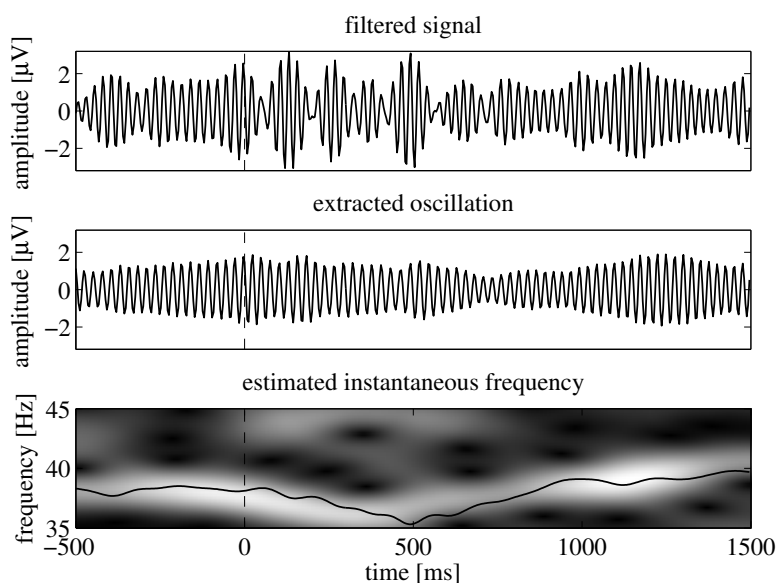


Figure 6.3: Pre-processing steps, (top) EEG signal filtered in the 35–45 Hz band, (middle) oscillation extracted with the SFT, and (bottom) its estimated instantaneous frequency. The estimated frequency is plotted on top of the short-time Fourier transform of the EEG signal. The vertical dashed lines denote stimulus onset.

power spectral density. This surrogate approach can help to highlight non-stationary effects, like stimulus-locked responses. This operation was repeated in order to obtain one surrogate signal for each EEG signal. Then, the same pre-processing was applied to the surrogate signals (fixed band-pass filtering and adaptive frequency tracking). This resulted in two datasets: a real dataset (corresponding to the real EEG signals) and a surrogate dataset. An example of surrogate signal is shown in Figure 6.4. One can observe that although the real and surrogate signals are different, they have the same amplitude spectrum.

6.2.3 Features

Once the pre-processing was applied to all EEG signals (real and surrogate ones), three different features were investigated. They were used for highlighting differences between real and surrogate datasets as well as between the two conditions used in the experiment (IC and NC). They also permitted to assess the usefulness of the SFT, as the same features were computed before and after frequency tracking. The features were computed on sliding windows of length 300 ms which offered a good tradeoff between temporal resolution and estimation accuracy. The time shift between successive windows was set to 10 ms. Sliding windows were used in order to visualize the evolution of the features over time. The complete procedure for extracting the features from the spatial-averaged EEG signals is depicted in Figure 6.1.

Then, statistical tests were performed in order to display significant differences between real and surrogate datasets or IC and NC conditions. However, before applying the tests, the features were transformed into approximately Gaussian variables whenever necessary. Finally, analysis of variance (ANOVA) [145] was performed with “subject” as a random effect and “dataset” or “condition” as a fixed effect for the 101 windows whose centers were located in the interval 0–1000 ms following stimulus presentation. Only the p-value for the fixed effect was considered. As many tests were performed for each feature which could lead to several type I errors [146],

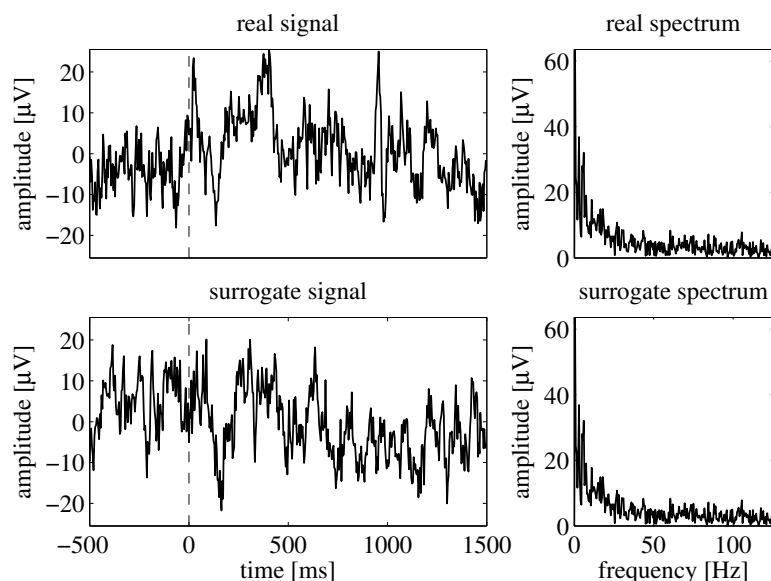


Figure 6.4: Original and surrogate EEG signals, (top left) real EEG signal, (bottom left) surrogate signal, (top right) amplitude spectrum of the real signal, (bottom right) amplitude spectrum of the surrogate signal. Vertical dashed lines denote stimulus onset.

an effect was declared significant only when the ANOVA yielded a p -value below 5% for some number of successive windows. Furthermore, permutation tests [147] were used to compute a lower bound for this number of successive significant windows in order to achieve a final p -value below 5%. These tests were performed by repeating 1000 times the ANOVA with randomly permuted dataset or condition memberships. In other words, the features computed on sliding windows were randomly reassigned to either of the datasets or conditions while keeping the true subjects' assignments and the natural temporal order of the windows so as to preserve the correlation structure. The p -values for all windows were then computed with the ANOVA, and the maximum number of successive significant windows ($p < 0.05$) was evaluated for the 1000 repetitions. Thus, the distribution of the maximum number of successive significant windows under the hypothesis of no difference between the datasets or conditions could be estimated. This led to an estimate of the probability of observing a number of successive significant windows equal to or greater than the one obtained with the true assignment of datasets or conditions when assuming no difference between these datasets or conditions. This probability estimate is in fact the p -value for the number of successive significant windows for the feature under investigation. In practice, a difference was declared significant only when this p -value was below 5% and the ANOVA rejected the null hypothesis for at least 4 successive windows. The latter condition ensured that the observed difference was not only anecdotal. This statistical analysis was applied to three features: the mean instantaneous frequency estimated by the SFT, the phase-amplitude and phase-phase couplings. All details regarding the computation of these features are provided in the following sections.

Mean frequency

The first feature considered in this study was the mean estimated instantaneous frequency, based on the estimate provided by the SFT. For this feature no transformation was needed for having approximately Gaussian-distributed variables due to the central limit effect [148]. Al-

though this study focused mostly on cross-frequency couplings, the mean frequency was also investigated as it was readily available thanks to the SFT. Also, this kind of feature was rarely analyzed in this context. It is important to notice that this feature could not be computed without adaptive frequency tracking.

Phase-amplitude couplings

Phase-amplitude couplings were measured with the phase locking value (PLV) [138]. The PLV is computed using the phase of the low frequency component $\phi_{lf}[n]$ and the phase of the amplitude of the high frequency component $\phi_{ahf}[n]$:

$$P_{PA} = \left| \mathbb{E} \left\{ e^{j(\phi_{lf}[n] - \phi_{ahf}[n])} \right\} \right|, \quad (6.1)$$

The phases and amplitudes were extracted with the DHT. The PLV takes a value of one for perfectly synchronized signals and zero when there is no synchronization. In practice, the expectation was replaced by the sample mean. The PLV values were transformed into approximately Gaussian random variables with an arcsine transform ($Z_{PA} = \arcsin(2P_{PA} - 1)$) before performing the statistical tests [149].

Phase-phase couplings

Phase-phase couplings were also measured with the PLV. However, this measure was slightly modified in order to take into account oscillatory components with different frequencies [137]. It is defined as follows for measuring phase-phase couplings:

$$P_{PP} = \left| \mathbb{E} \left\{ e^{j(a\phi_{lf}[n] - b\phi_{hf}[n])} \right\} \right|, \quad (6.2)$$

where a and b are coupling coefficients. The following values for $a:b$ were considered: 4:3, 3:2, 2:1, 3:1, 4:1, 5:1, 6:1, 7:1, 8:1 and 9:1. Only the PLVs for which the frequency bands and coupling coefficients made sense were computed (i.e. bands that overlapped once multiplied by the corresponding coupling coefficients). For instance, when measuring the phase-phase couplings between 1–4 Hz and 35–45 Hz components, the coupling coefficients 2:1 were discarded. Indeed, multiplying the limits of each band with the corresponding coefficient yields 2–8 Hz and 35–45 Hz which do not overlap. As for the phase-amplitude couplings, the expectation was replaced by the sample mean, and the same arcsine transform was applied before performing the statistical analyses ($Z_{PP} = \arcsin(2P_{PP} - 1)$).

6.2.4 Synthetic Signals

The usefulness of the SFT for measuring cross-frequency couplings was also evaluated with synthetic signals and Monte Carlo simulations. Two cases were considered: a basic case with sinusoids embedded in additive noise and a more complex case in which synthetic signals were generated in order to mimic a real EEG signal. In the first case, the goal was to measure the phase-phase couplings with the PLV between two simple signals. The input signals were defined as two sinusoids at normalized frequencies 0.05 and 0.35 with uniformly distributed phases embedded in independent white Gaussian noises:

$$x_1[n] = \sin(2\pi 0.05n + \theta_1) + v_1[n], \quad n = 0, 1, \dots, 1074, \quad (6.3a)$$

$$x_2[n] = \frac{1}{3} \sin(2\pi 0.35n + \theta_2) + v_2[n], \quad n = 0, 1, \dots, 1074, \quad (6.3b)$$

where θ_1 and θ_2 were random phase offsets, and $v_1[n]$ and $v_2[n]$ were the additive noise processes. The 1/3 factor was used to take into account the power decrease in higher frequencies in EEG

data. Then, the SFT was applied for extracting the oscillatory component in each signal. The parameters of the adaptive algorithm were set to $\beta = 0.975$ and $\delta = 0.95$. The phases of input and output signals were extracted with the DHT. The first and last 500 samples were then discarded in order to avoid any border effect. This yielded 75-samples phase signals whose length corresponded to the one of the 300-ms windows used in the EEG analysis. Finally, the PLV was computed with coupling coefficients set to 7:1 for the phase signals obtained with and without frequency tracking. The PLV mean and standard deviation were estimated with 10,000 Monte Carlo simulations for the two approaches. Furthermore, this procedure was repeated for SNR values ranging from 0 to 20 dB in 1 dB steps. It is important to note that without noise the PLV for these signals should be equal to one.

In the second case, we generated two 500-samples signals mimicking the outputs of the band-pass filters used when analyzing the real EEG data. The sampling frequency for these two synthetic signals was set to 250 Hz which corresponded to a duration of 2000 ms. The first signal was defined as

$$x_1[n] = 5A[n] \sin\left(2\pi \frac{5.5}{250}n + \theta_1\right) + v_1[n], \quad n = 0, 1, \dots, 499, \quad (6.4a)$$

and the second one as

$$x_2[n] = 2A[n] \sin\left(2\pi \frac{44}{250}n + \theta_2\right) + u[n] + v_2[n], \quad n = 0, 1, \dots, 499, \quad (6.4b)$$

where θ_1 and θ_2 were random phase terms, $v_1[n]$ and $v_2[n]$ were additive white Gaussian noises with variances 25 and 4 respectively, $A[n]$ was the time-varying amplitude of the sinusoids and $u[n]$ was a transient periodic interference at 36 Hz. They were set to

$$A[n] = \begin{cases} 0.2 & \text{for } 0 \leq n < 125, \\ 0.2 + 0.8(n - 125)/25 & \text{for } 125 \leq n < 150, \\ 1 & \text{for } 150 \leq n < 350, \\ 1 - 0.8(n - 350)/25 & \text{for } 350 \leq n < 375, \\ 0.2 & \text{for } 375 \leq n < 500, \end{cases} \quad (6.5)$$

and

$$u[n] = \begin{cases} 0 & \text{for } 0 \leq n < 125, \\ 3g[n - 125] \sin(2\pi(36/250)n + \theta_3) & \text{for } 125 \leq n < 375, \\ 0 & \text{for } 375 \leq n < 500, \end{cases} \quad (6.6)$$

where $g[n]$ is a 250-samples Hann window [31] and θ_3 is a random phase offset. The signals $x_1[n]$, $x_2[n]$, $A[n]$ and $u[n]$ are illustrated in Figure 6.5 for one realization of the random phases and noises. The first signal was then filtered in the 4–8 Hz band and the second one in the 35–45 Hz band with the same fixed band-pass filters used before. All the parameters were chosen in order to generate synthetic signals inspired by a real EEG signal with stable oscillatory components at 5.5 Hz and 44 Hz and a short periodic interference at 36 Hz. The SFT was applied to both signals for extracting the main periodic components with the same parameters and mirroring procedure as for the real EEG data. The phase-phase couplings were then measured by computing the PLV with coupling coefficients 8:1 over sliding windows of length 300 ms shifted by 10 ms. The results obtained with and without tracking were averaged over 10,000 realizations of the random phases and noises. In this synthetic example, there should be an increase in coupling strength when n is between 150 and 350 samples (600–1400 ms).

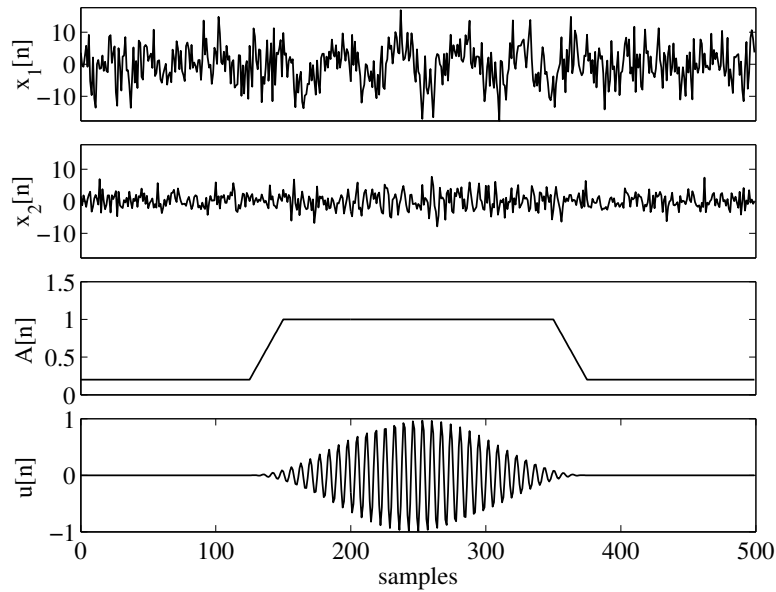


Figure 6.5: Examples of the synthetic signals used for assessing the usefulness of the SFT, from top to bottom: $x_1[n]$ (6.4a), $x_2[n]$ (6.4b), $A[n]$ (6.5) and $u[n]$ (6.6).

6.3 Results

6.3.1 Synthetic Signals

Before presenting the results obtained with the signals recorded during the IC experiment, the outcomes of the Monte Carlo simulations with synthetic signals are presented. The mean PLVs with error bars obtained with the first set of signals (6.3) are shown in Figure 6.6 for all tested SNR values. Without noise the PLV should be equal to one in this scenario as the two oscillatory components were perfectly synchronized with 7:1 coefficients. This was indeed the case for very low noise levels. However, without adaptive tracking, the mean PLV quickly decreased as the noise variance increased. This decrease was quite severe even for moderate noise levels. By contrast, the SFT led to PLV values that were much more resilient to noise, at the cost of increased estimation variance however. Nevertheless, the SFT increased the overall performance of the PLV for measuring phase-phase couplings with these synthetic signals. Indeed, although the PLV variance was higher with tracking, the mean PLV obtained without tracking reached its minimal value for SNR values below 5 dB. In the second case, where synthetic signals were generated so as to imitate real EEG data (6.4), the SFT also proved to be helpful for measuring phase-phase couplings. These signals contained two perfectly synchronized sinusoids with 8:1 coefficients embedded in noise. An interfering periodic component active during a short duration was also present in the high-frequency signal. The two sinusoids had time-varying amplitudes that reached their maximal values in the interval 600–1400 ms. And thus, the PLV computed over 300-ms sliding windows shifted by 10 ms should increase in this interval. The results averaged over 10,000 Monte Carlo simulations are shown in Figure 6.7 with and without adaptive frequency tracking. Without the SFT, the PLV increased in the beginning of the interval as expected, but it was then completely disrupted by the interfering oscillation. In the end of the interval, it increased again as the amplitude of the interference dropped. On the contrary, the PLV values obtained with the SFT were higher during the whole duration of the interval, except

for an adaptation delay (~150 ms). Therefore, with these synthetic signals, meaningful phase information could be extracted by means of the SFT, which led to robust PLV values. In particular, the first example showed its tolerance to broad-band noise, while the second one showed its resilience to interfering oscillations. This concludes the analysis of synthetic signals, and thus the following sections focus on real EEG data.

6.3.2 Mean Frequency

The comparisons between the real and surrogate datasets yielded significant differences in terms of mean frequency for the 1–4 Hz component in the interval 180–380 ms (21 successive windows, permutation test: $p < 0.001$) and the 4–8 Hz component in the interval 200–380 ms (19 successive windows, permutation test: $p < 0.001$). In both cases, the estimated instantaneous frequency was higher for the real data than for the surrogate ones. This is illustrated in Figure 6.8. It seems that the mean frequency of the main oscillatory component in the band 1–4 Hz increased smoothly following stimulus presentation for the two datasets. However, this frequency increase was more important in the real data. The phenomenon was slightly different for the 4–8 Hz band. Indeed, in this case, the main frequency remained almost constant for the surrogate dataset. By contrast, an increase in mean frequency after stimulus presentation for the real dataset caused the significant difference. There was no other significant difference between the two datasets for this feature.

As significant differences in mean frequency were observed between the real and surrogate datasets, the IC and NC conditions were also compared with respect to this feature. There were significant differences between the two conditions for the instantaneous frequency of the components in the 4–8 Hz and 45–55 Hz bands. These two comparisons are shown in Figure 6.9. The frequency was significantly higher for IC than for NC for the 4–8 Hz component in the interval 230–610 ms (39 successive windows, permutation test: $p < 0.001$). The frequency increase following stimulus presentation was more pronounced and lasted longer for IC than for NC. By contrast, IC yielded a lower mean frequency than NC for the 45–55 Hz component. This difference was significant in the interval 680–920 ms (25 successive windows, permutation test: $p < 0.05$). In this case, the frequency decreased more than 500 ms after stimulus for IC while it increased for NC. No other significant difference between the conditions was observed in terms of frequency.

6.3.3 Phase-amplitude Couplings

Several combinations of components yielded significant differences after stimulus onset in phase-amplitude couplings when comparing real and surrogate datasets. Most of the significant results were obtained when the 4–8 Hz component was involved. A significant decrease in coupling strength was observed after stimulus presentation for the real dataset compared to the surrogate one when the 4–8 Hz component was considered as the low frequency component. By contrast, an increase was obtained with the components from the 1–4 Hz and 4–8 Hz bands. All these results involving the 4–8 Hz component are summarized in Figure 6.14. Furthermore, two examples of comparisons between the real and surrogate datasets in terms of phase-amplitude couplings are shown in Figure 6.10 for the combinations 1–4 Hz and 4–8 Hz as well as 4–8 Hz and 35–45 Hz. Other significant results were obtained with the following pairs of components: 1–4 Hz and 8–12 Hz in the interval 170–320 ms (real > surrogate, 16 successive windows, permutation test: $p < 0.001$), 1–4 Hz and 15–25 Hz in the interval 190–420 ms (real > surrogate, 24 successive windows, permutation test: $p < 0.001$), 1–4 Hz and 55–65 Hz in the interval 220–340 ms (real > surrogate, 13 successive windows, permutation test: $p < 0.01$), 8–12 Hz and 15–25 Hz in the interval 90–280 ms (real < surrogate, 20 successive windows, permutation

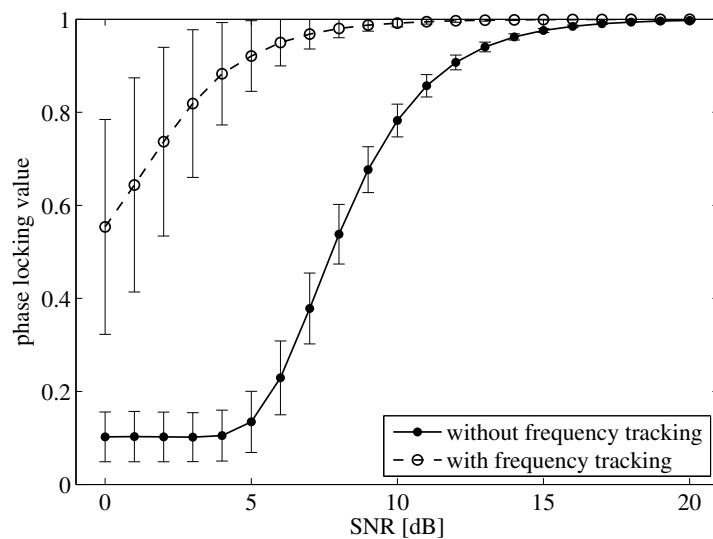


Figure 6.6: Mean PLV values with corresponding error bars obtained with and without adaptive frequency tracking for SNR values ranging from 0 to 20 dB in 1 dB steps.

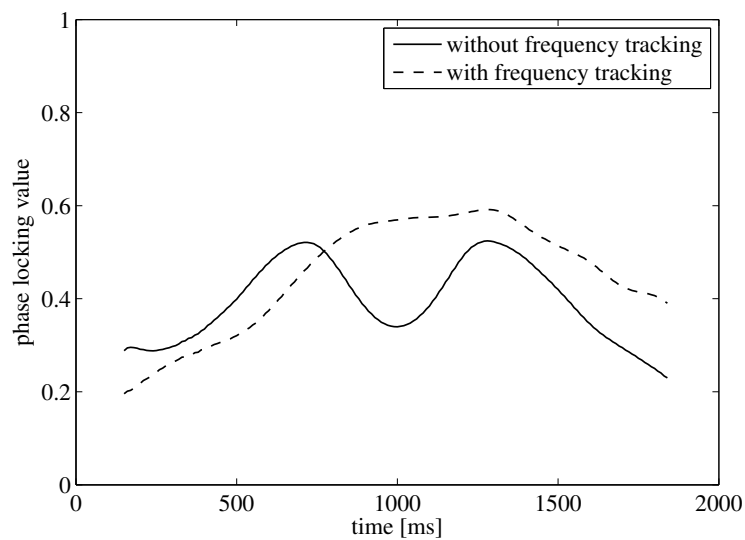


Figure 6.7: Mean PLV values measuring the phase-phase couplings over 300-ms sliding windows for a synthetic EEG example.

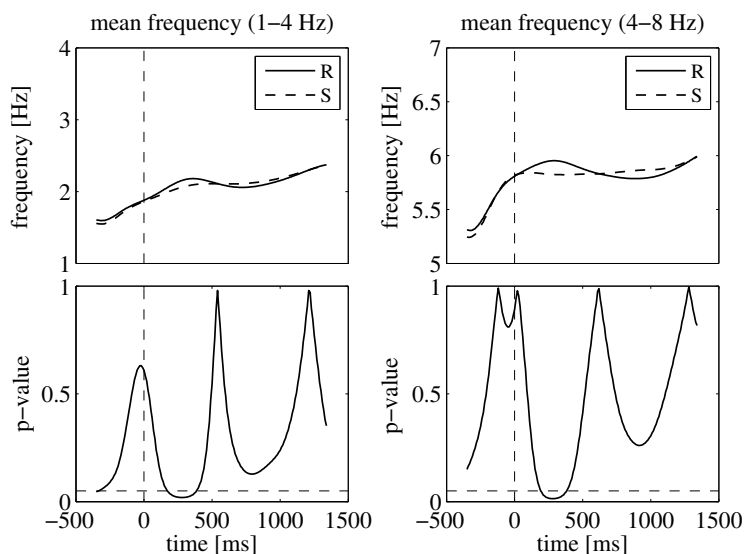


Figure 6.8: Comparisons between real (R) and surrogate (S) datasets for the mean estimated frequency. Mean frequencies for the (top left) 1–4 Hz and (top right) 4–8 Hz components, and ANOVA p-values for the comparison between the two datasets for the (bottom left) 1–4 Hz and (bottom right) 4–8 Hz components. Vertical dashed lines denote stimulus onset and horizontal ones denote the 5% significance level.

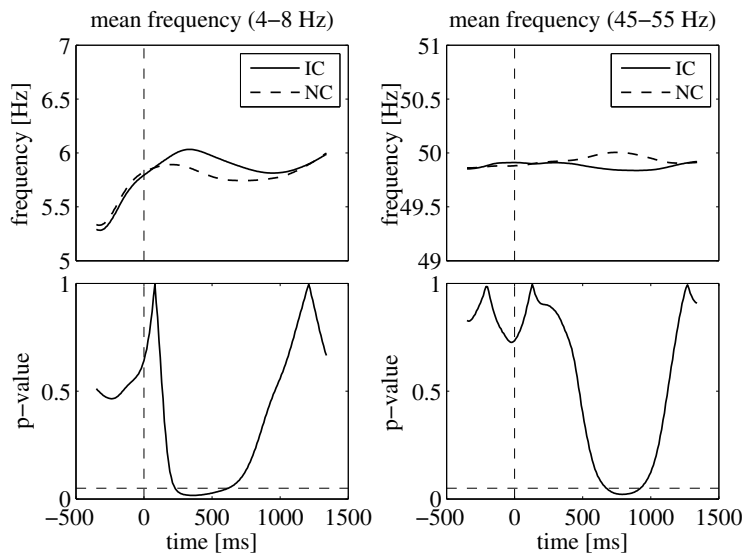


Figure 6.9: Comparisons between IC and NC conditions for the mean estimated frequency. Mean frequencies for the (top left) 4–8 Hz and (top right) 45–55 Hz components, and ANOVA p-values for the comparison between the two conditions for the (bottom left) 4–8 Hz and (bottom right) 45–55 Hz components. Vertical dashed lines denote stimulus onset and horizontal ones denote the 5% significance level.

test: $p < 0.001$), and 8–12 Hz and 35–45 Hz in the interval 210–240 ms (real < surrogate, 4 successive windows, permutation test: $p < 0.05$).

The strength of phase-amplitude couplings for the IC and NC conditions was also compared. All significant differences were found when the 4–8 Hz component was involved. In fact, when the 4–8 Hz component was considered as the low frequency component, the coupling strength was smaller for IC than for NC. It was the inverse when the 4–8 Hz component was considered as the high frequency component (this was only the case for the combination 1–4 Hz and 4–8 Hz). An example of this phenomenon is illustrated in Figure 6.11. It can be seen that the stimulus caused a change in phase-amplitude coupling strength (either an increase or a decrease) and that this change was always more pronounced for the IC than for the NC condition. All the results of the differences between the two conditions for the phase-amplitude couplings involving the 4–8 Hz component are illustrated in Figure 6.15.

6.3.4 Phase-phase Couplings

The pattern of results was more complex for the phase-phase couplings. Different results were obtained for different coupling coefficients. However, similarly to the phase-amplitude couplings, most of the significant differences were observed when the 4–8 Hz component was considered either as the low- or high-frequency component. When comparing real and surrogate datasets, the dataset yielding higher coupling strength varied depending on the ratio of coupling coefficients. The phase-phase coupling between the 1–4 Hz and 4–8 Hz components were lower for the real dataset than for the surrogate one when a low coefficient ratio was used (4:3 and 3:2). On the contrary, when this ratio increased (from 4:1), the coupling strength was higher for the real signals. For coefficients 2:1, the coupling strength was higher for the real data while it was lower for coefficients 3:1. When the 4–8 Hz component was considered as the low frequency component, the coupling strength was higher for low ratios of coefficients. But, as before, the opposite result was observed for higher ratios. An example of this inversion phenomenon is shown in Figure 6.12 for the 4–8 Hz and 35–45 Hz components with coefficient pairs set to 6:1 and 9:1. One can observe that these significant differences were caused by sharp changes (either increase or decrease) in the PLV. All the results of the comparisons between real and surrogate datasets for phase-phase couplings for the 4–8 Hz component are summarized in Figure 6.14. The permutation tests also identified significant differences between the two datasets for several other combinations of bands. The combinations of bands for which the coupling strength was significantly higher for the real data were the 1–4 Hz and 8–12 Hz bands with coefficients 3:1, 7:1, 8:1 and 9:1, the 1–4 Hz and 15–25 Hz bands with coefficients 6:1, the 1–4 Hz and 35–45 Hz bands with coefficients 9:1, the 8–12 Hz and 15–25 Hz bands with coefficients 2:1 and 3:1, the 8–12 Hz and 35–45 Hz bands with coefficients 4:1, the 8–12 Hz and 45–55 Hz bands with coefficients 6:1, the 8–12 Hz and 55–65 Hz bands with coefficients 8:1, the 8–12 Hz and 65–75 Hz bands with coefficients 9:1, the 15–25 Hz and 35–45 Hz bands with coefficients 2:1. On the other hand, the combinations of bands showing significantly higher coupling strength for the surrogate datasets were the 1–4 Hz and 8–12 Hz band with coefficients 4:1 and 5:1, the 8–12 Hz and 35–45 Hz bands with coefficients 3:1, the 8–12 Hz and 45–55 Hz bands with coefficients 4:1 and 5:1, the 8–12 Hz and 55–65 Hz bands with coefficients 6:1, the 8–12 Hz and 65–75 Hz bands with coefficients 6:1 and 7:1. Consequently, the inversion phenomenon was also observed for other combinations of bands, in particular the ones involving the 8–12 Hz band.

A very similar inversion phenomenon depending on coupling coefficients occurred when comparing the IC and NC conditions for the phase-phase couplings. When measuring phase-phase couplings between the 1–4 Hz and 4–8 Hz components, higher coupling strength was observed for NC with low ratios of coefficients (4:3 and 3:2). While, for the coefficient 3:1, IC yielded higher coupling strength. For larger ratios, no clear differences were found for these two bands. When the 4–8 Hz component was considered as the low frequency component, the IC

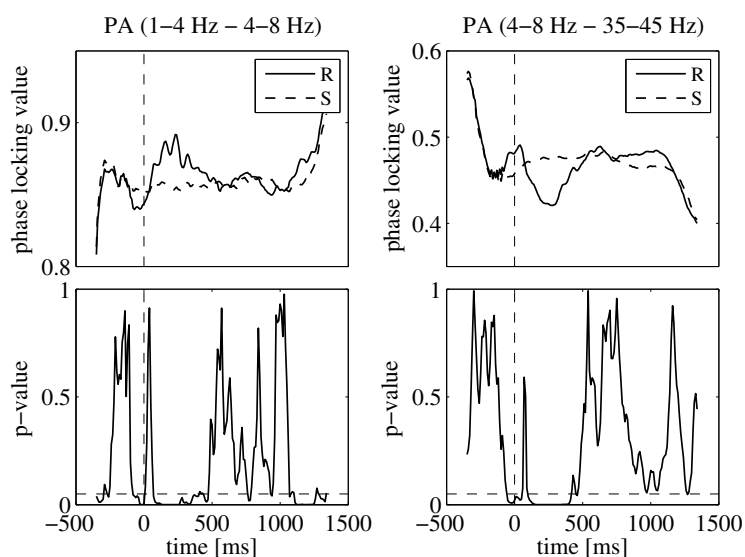


Figure 6.10: Comparisons between real (R) and surrogate (S) datasets for the phase-amplitude (PA) couplings measured using the PLV. Mean PLVs for (top left) the 1–4 Hz and 4–8 Hz components and (top right) the 4–8 Hz and 35–45 Hz components, and ANOVA p-values for the comparison between the two datasets for (bottom left) the 1–4 Hz and 4–8 Hz components and (top right) the 4–8 Hz and 35–45 Hz components. Vertical dashed lines denote stimulus onset and horizontal ones denote the 5% significance level.

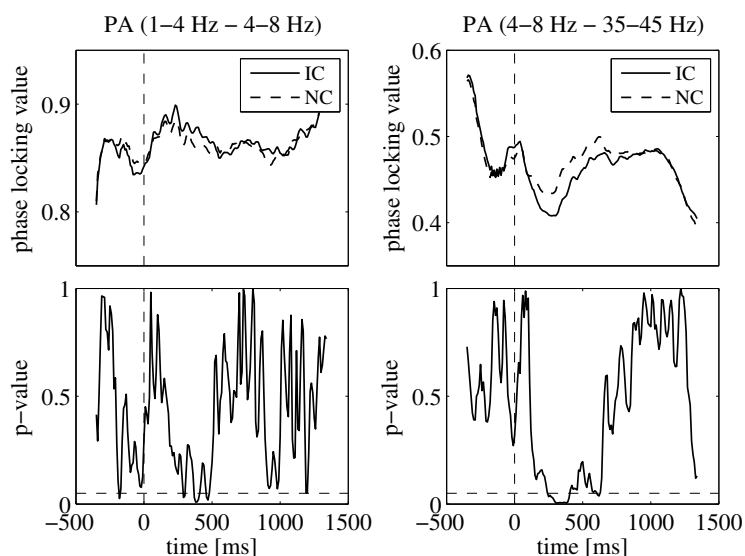


Figure 6.11: Comparisons between IC and NC conditions for the phase-amplitude (PA) couplings measured using the PLV. Mean PLVs for (top left) the 1–4 Hz and 4–8 Hz components and (top right) the 4–8 Hz and 35–45 Hz components, and ANOVA p-values for the comparison between the two conditions for (bottom left) the 1–4 Hz and 4–8 Hz components and (top right) the 4–8 Hz and 35–45 Hz components. Vertical dashed lines denote stimulus onset and horizontal ones denote the 5% significance level.

condition led to higher coupling strength compared to the NC condition for low ratios of coefficient pairs. However, as the ratio increased, the condition yielding the higher coupling strength changed to NC. Figure 6.13 shows an example of this change for the 4–8 Hz and 8–12 Hz components with coefficients set to 3:2 and 2:1. This phenomenon was observed for various combinations of frequency bands and coefficient pairs. All the results obtained with the 4–8 Hz component are reported in Figure 6.15.

6.3.5 Advantages of Adaptive Frequency Tracking

For assessing the usefulness of the SFT for real EEG data, the two conditions IC and NC were compared when the features were computed without adaptive frequency tracking. In other words, the features were also computed using the output signals of the predefined band-pass filters as shown in Figure 6.1. Obviously, the mean frequency could not be estimated without frequency tracking as it is specifically an output of the SFT. Nevertheless, the phase-amplitude and phase-phase couplings were measured and the IC and NC conditions were compared with the same test procedure as before. The focus was put on cross-frequency couplings involving the 4–8 Hz component as they yielded all the significant differences when the SFT was applied. Similarly to Figure 6.15, the results obtained in this case are shown in Figure 6.16. Comparing the two figures, one can notice that the results obtained with and without adaptive frequency tracking are very similar and that there is no conflict. However, a more detailed investigation revealed that the differences between IC and NC in terms of coupling strength were, in most cases, more clearly highlighted when using the SFT. In fact, when investigating the differences in cross-frequency couplings involving the 4–8 Hz component, a greater number of successive significant windows was obtained for only seven cases without tracking, including two cases where no significant difference was observed after applying the SFT. For all other cases, the proposed adaptive scheme performed as well as or (more frequently) better than the traditional approach, in terms of number of successive significant windows. It also led to the detection of significant differences between IC and NC which remained unnoticed without tracking in five cases. The usefulness of the SFT is particularly apparent for phase-amplitude couplings. However, it is worth mentioning that the adaptation process in this algorithm introduces a delay. This caused the intervals of significant differences to be shifted in time compared the ones obtained without tracking. A coarse method for comparing the results obtained with and without the SFT is to count the number of successive significant windows in Figures 6.15 and 6.16 for phase-amplitude and phase-phase couplings. Thus, for the phase-amplitude couplings, this led to 102 and 29 windows with and without frequency tracking. These values were 272 and 208 for phase-phase couplings. As mentioned previously, when Figures 6.15 and 6.16 are put side-by-side, the significant intervals obtained with the SFT were delayed compared to those obtained without tracking as the adaptation process of the proposed algorithm is not instantaneous. This delay could be quite large depending on the dynamics of the signals under study. Nevertheless, in most cases, it remained reasonable and the intervals overlapped. However, there were a few cases where the delay was quite important (> 200 ms). For instance, the significant interval obtained when comparing the phase-phase couplings with 4:3 coefficients between the 1–4 Hz and 4–8 Hz components occurred 300 ms later with frequency tracking. These long delays, while unusual, could be explained by large non-stationary dynamics in the input signals of the SFT.

For the sake of completeness, the extents and significance levels of the intervals shown in Figures 6.14, 6.15 and 6.16 are also reported in numeric format at the end of this chapter in Tables 6.1, 6.2 and 6.3.

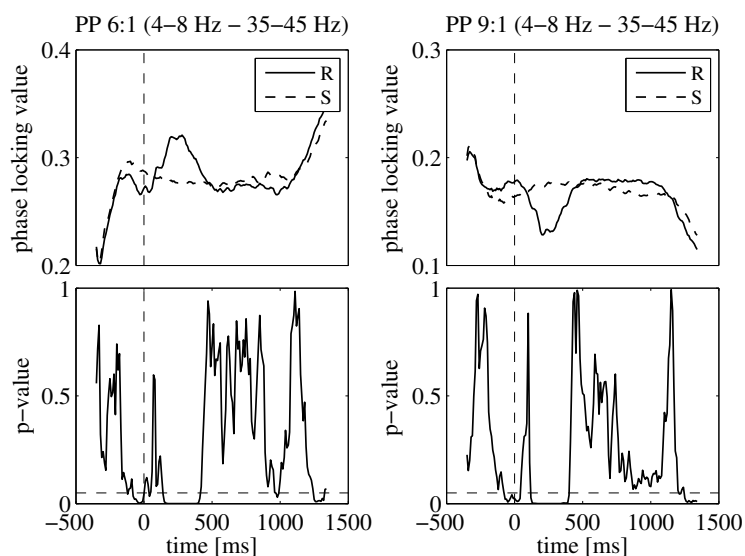


Figure 6.12: Comparisons between real (R) and surrogate (S) datasets for the phase-phase (PP) couplings measured using the PLV between the 4–8 Hz and 35–45 Hz components. Mean PLVs with coupling coefficients set to (top left) 6:1 and (top right) 9:1, and ANOVA p-values for the comparison between the two datasets with coefficients set to (bottom left) 6:1 and (bottom right) 9:1. Vertical dashed lines denote stimulus onset and horizontal ones denote the 5% significance level.

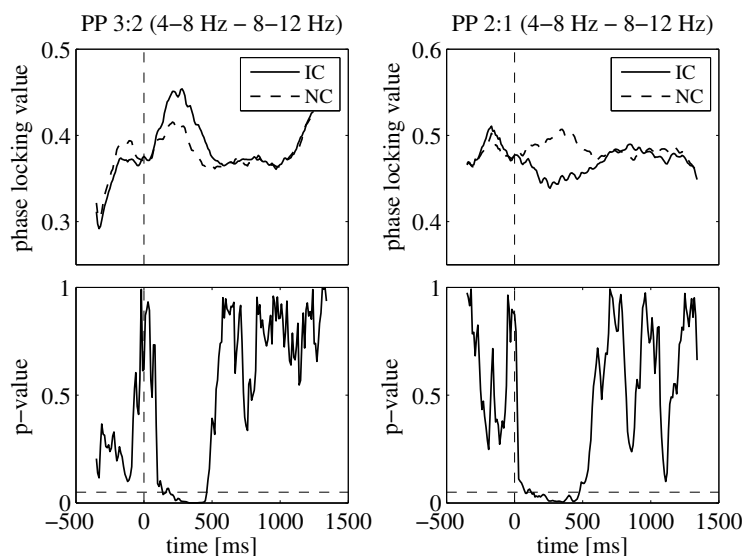


Figure 6.13: Comparisons between IC and NC conditions for the phase-phase (PP) couplings measured using the PLV between the 4–8 Hz and 8–12 Hz components. Mean PLVs with coupling coefficients set to (top left) 3:2 and (top right) 2:1, and ANOVA p-values for the comparison between the two conditions with coefficients set to (bottom left) 3:2 and (bottom right) 2:1. Vertical dashed lines denote stimulus onset and horizontal ones denote the 5% significance level.

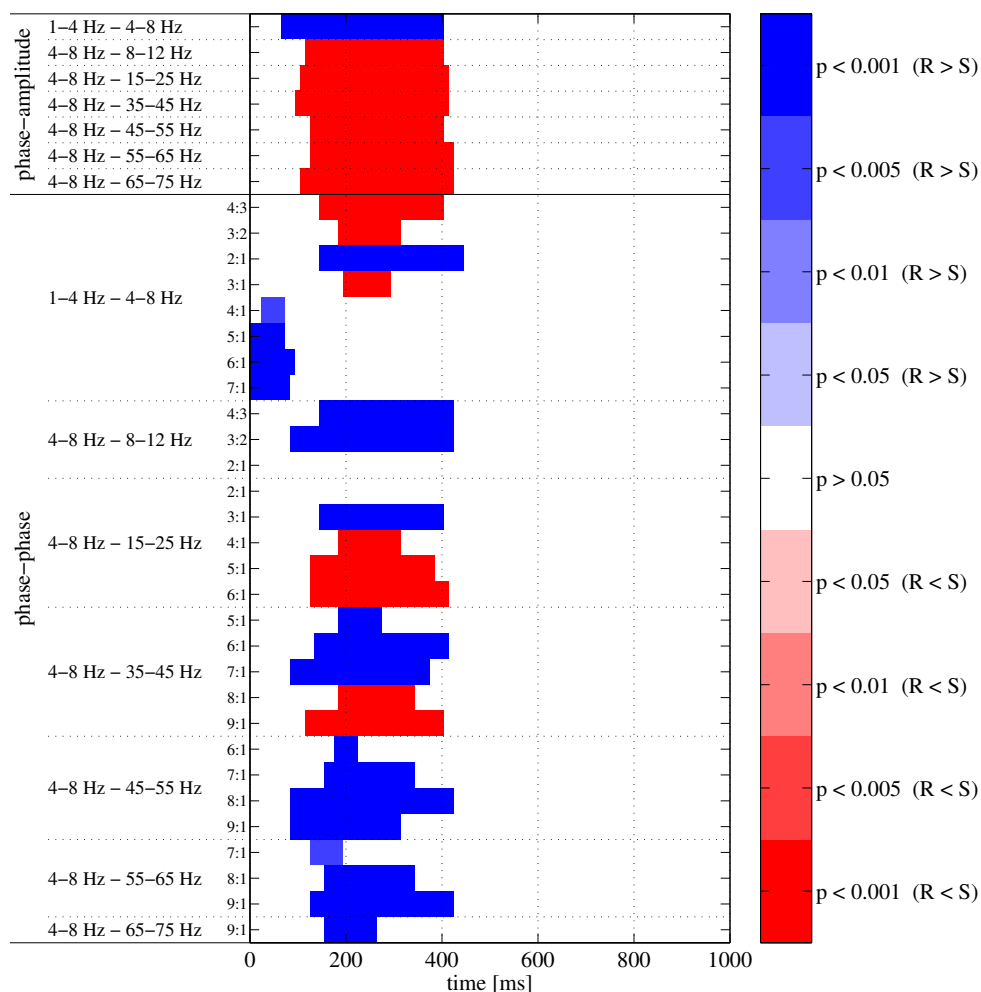


Figure 6.14: Comparisons between real (R) and surrogate (S) datasets for phase-amplitude (top rows) and phase-phase (bottom rows) couplings measured with the SFT. Significant intervals are shown in blue (respectively in red) when the coupling strength was higher for the real (respectively surrogate) dataset. Color intensity denotes the significance level of the corresponding permutation test.

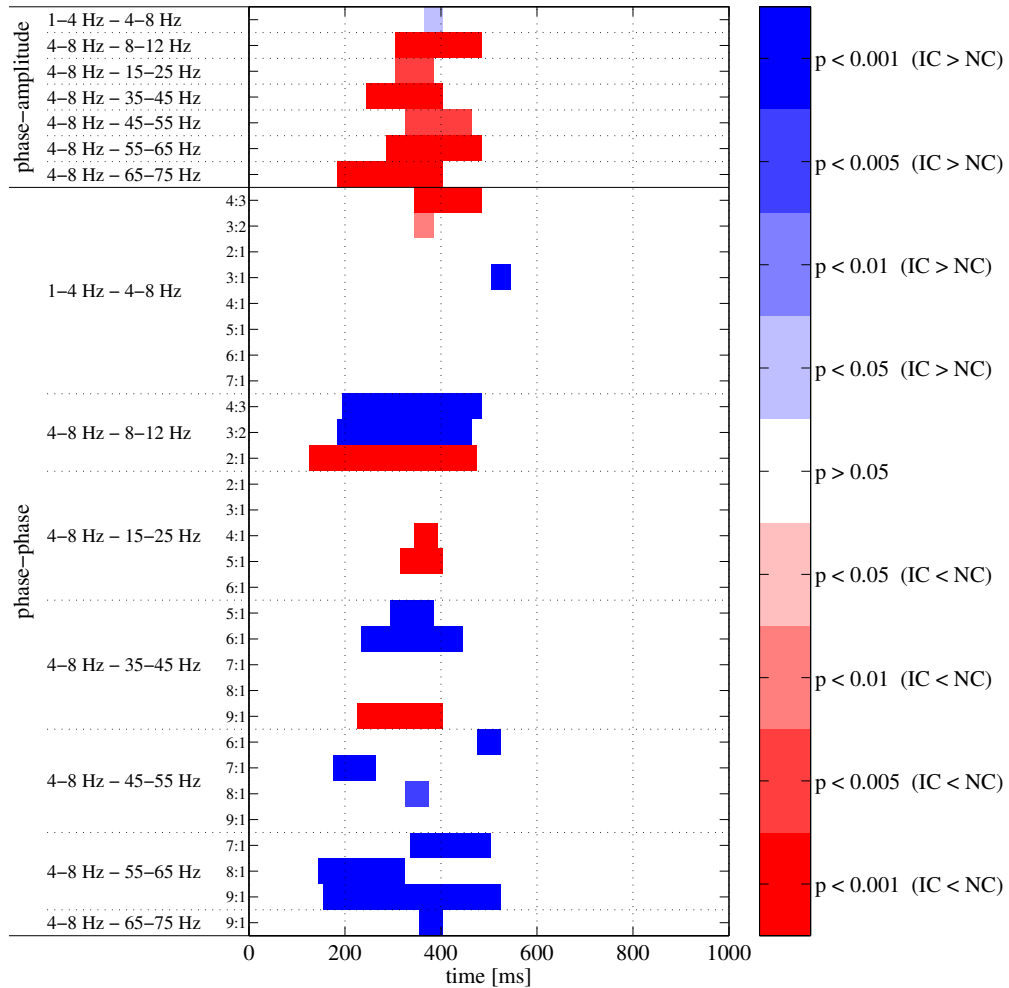


Figure 6.15: Comparisons between IC and NC conditions for phase-amplitude (top rows) and phase-phase (bottom rows) couplings measured with the SFT. Significant intervals are shown in blue (respectively in red) when the coupling strength was higher for the IC (respectively NC) condition. Color intensity denotes the significance level of the corresponding permutation test.

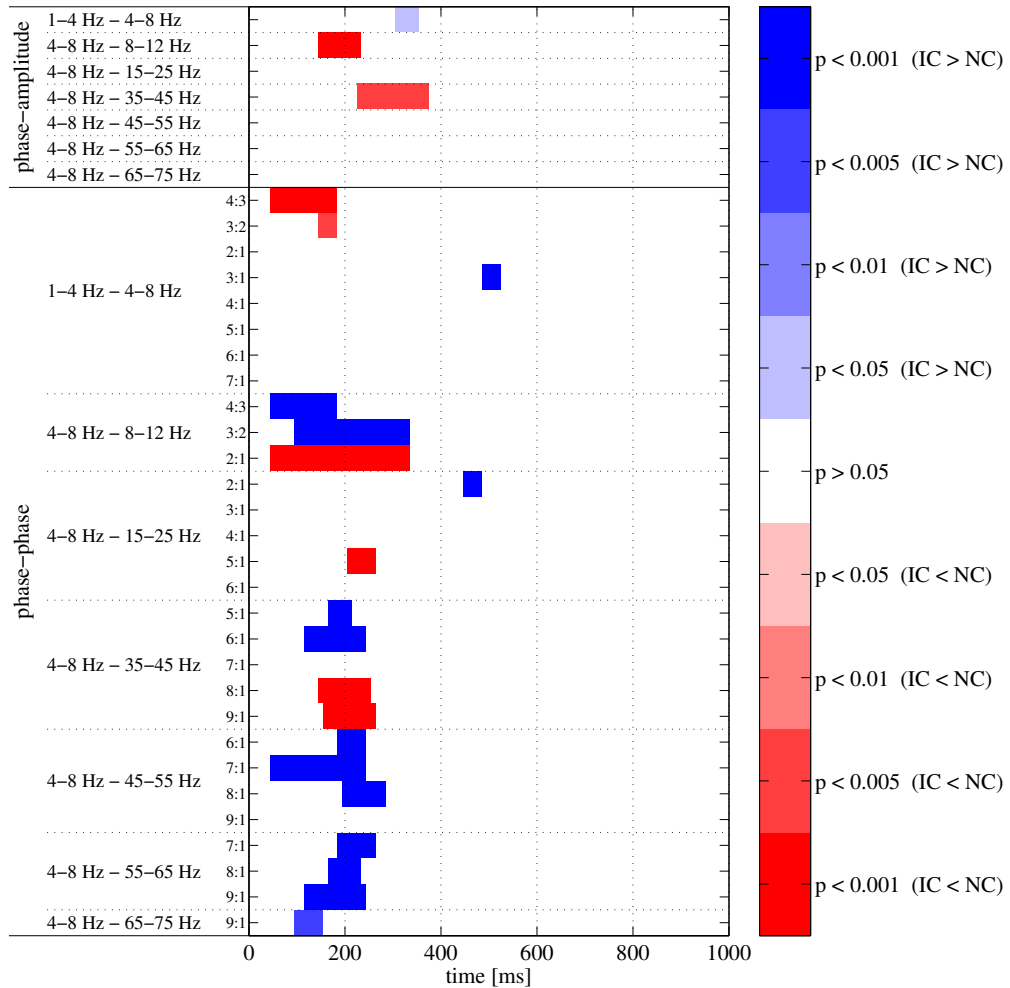


Figure 6.16: Comparisons between IC and NC conditions for phase-amplitude (top rows) and phase-phase (bottom rows) couplings measured without the SFT. Significant intervals are shown in blue (respectively in red) when the coupling strength was higher for the IC (respectively NC) condition. Color intensity denotes the significance level of the corresponding permutation test.

6.4 Discussion

Advances in analysis methods have revealed the importance of neuronal oscillations in brain activity and function. Recent studies have highlighted that the top-down control of perception and brain responses is supported to a large extent by oscillatory activity [70]. Consequently, these oscillatory components are now considered as highly efficient information-rich signals in the field of neuroscience. Furthermore, the coupling mechanisms occurring across frequency bands have been the focus of several recent studies [119, 150–152]. Collectively, these findings prompted the development of an adaptive frequency tracking scheme, the SFT, for analyzing EEG data in more detail. Specifically, this algorithm was designed to maximize the oscillatory behavior at the output, which is very important for extracting proper phase information, which, in turn, can be used to measure cross-frequency couplings.

The advantages of the SFT for measuring cross-frequency couplings were evaluated with synthetic signals and real EEG data recorded during an IC experiment. First, the synthetic signals in conjunction with Monte Carlo simulations highlighted two desirable features of the proposed algorithm. In the first case, it was shown to be resilient to broad-band noise as the PLV decrease remained limited in high noise levels (Figure 6.6). In the second case, synthetic signals imitating real EEG recordings were generated in order to check that the SFT could cope well with interfering oscillatory components (Figure 6.7). Therefore, these numerical simulations illustrated two advantages of the adaptive scheme (resilience to broad-band noise and oscillatory interference) compared to classical filter-bank approaches. These advantages were confirmed when the SFT was applied to real EEG signals for extracting the temporal evolution of differences between the IC and NC conditions in terms of phase-amplitude and phase-phase couplings. The number of successive significant windows was larger with tracking than without for almost all combinations of bands. The advantages of the adaptive algorithm were particularly apparent for phase-amplitude couplings. Furthermore, although two significant differences of phase-phase couplings were only detected without the SFT, it led to the detection of five such differences that remained unnoticed with traditional band-pass filtering. And the lengths of the significant intervals were longer with adaptive frequency tracking in most cases. Thus, adaptive frequency tracking could improve the measurements of cross-frequency couplings through precise extraction of neuronal oscillations. Moreover, as the SFT also provides an estimate of the instantaneous frequency of the extracted component, significant changes in frequency could be observed for a few of the bands under study, both when comparing the real and surrogate datasets and the two conditions.

When considering more closely the outcomes of the comparisons between datasets and conditions (Figures 6.14 and 6.15), a complex pattern of results was highlighted by the SFT. Nonetheless, a few important observations can be pointed out. First, the dataset or condition yielding the highest coupling strength depended on the combination of bands. And, for phase-phase couplings, it also depended on the coupling coefficients, or more specifically on the coefficient ratio. Second, when comparing real and surrogate datasets, the significant differences were in most cases due to changes in coupling strength, either decreases or increases, for the real signals while it remained more or less constant for the surrogate ones. Since the surrogate data were generated so as to be stationary, it was expected. A similar phenomenon was observed when comparing IC and NC conditions. However, usually both conditions elicited a change in coupling strength in the same direction. Nevertheless, this change was typically more pronounced for IC. This seems to indicate that the processing of such contours requires more changes in terms of cross-frequency couplings, but clearly more investigations are needed to confirm this observation. Regarding the outcomes of the comparisons between the two conditions, it is important to mention that the increase of the instantaneous frequency observed of the 4–8 Hz component observed during IC processing (Figure 6.9) was too weak to account for the inversion of the differences in phase-phase couplings depending on the coefficient ratio. Clearly, more investigations are re-

quired to perfectly understand the role of cross-frequency couplings. Nonetheless, the coupling mechanisms reported in this chapter may link the responses to visual stimuli observed in the lower frequencies [85, 131–133, 153, 154] to the ones observed in the higher frequencies [155], alongside the results of numerous studies about cross-frequency couplings [119, 135, 156, 157].

Some limitations concerning this study and the proposed algorithm are worth discussing. First, the adaptation process in the SFT is not instantaneous, and consequently the estimated frequency suffers some delay. Thus, the time-varying band-pass filter used for extraction needs some time to center on the tracked periodic component. This delay not only depends on the SFT parameters, but also on the dynamics of the signal of interest. Indeed, the adaptation is slower in highly non-stationary environments. The delay introduced by the algorithm is clearly visible when comparing the significant intervals for cross-frequency couplings measured with and without adaptive frequency tracking (Figures 6.15 and 6.16). Future developments of the SFT will address this issue. In particular, an approach for compensating the delay [80] may prove useful in the framework of EEG processing. This study also focused on only a small cluster of surface electrodes chosen on the basis of the results of a previous investigation [85], and therefore only local information regarding the cross-frequency couplings were obtained. Ideally, the same analysis procedure should be repeated for all available electrodes in future studies. Furthermore, adaptive frequency tracking could also be applied to intracranial EEG signals or to signals computed through inverse solution (e.g. [158]). In particular, the second type of signals can also be used to measure couplings not only across frequency bands but also across different brain areas (e.g. [159–161]). However, as the number of signals to analyze increases the processing time may become prohibitive. The computation load comes mainly from the statistical analysis based on permutation tests as the other operations such as filtering and tracking are fairly time-efficient. The coupling analysis performed in this study also raises another important question concerning the direction of the cross-frequency interactions: are the low-frequency oscillations controlling the high-frequency ones or is it the inverse [162]? This issue can be investigated with causality measures, however they still have some drawbacks that may render them inefficient in this case.

	1–4 Hz 4–8 Hz	4–8 Hz 8–12 Hz	4–8 Hz 15–25 Hz	4–8 Hz 35–45 Hz	4–8 Hz 45–55 Hz	4–8 Hz 55–65 Hz	4–8 Hz 65–75 Hz
PA	R > S 34: 70–400 $p < 0.001$	R < S 29: 120–400 $p < 0.001$	R < S 31: 110–410 $p < 0.001$	R < S 32: 100–410 $p < 0.001$	R < S 28: 130–400 $p < 0.001$	R < S 30: 130–420 $p < 0.001$	R < S 32: 110–420 $p < 0.001$
PP 4:3	R < S 26: 150–400 $p < 0.001$	R > S 28: 150–420 $p < 0.001$					
PP 3:2	R < S 13: 190–310 $p < 0.001$	R > S 34: 90–420 $p < 0.001$					
PP 2:1	R > S 30: 150–440 $p < 0.001$	R = S	R = S				
PP 3:1	R < S 10: 200–290 $p < 0.001$		R > S 26: 150–400 $p < 0.001$				
PP 4:1	R > S 5: 30–70 $p < 0.005$		R < S 13: 190–310 $p < 0.001$				
PP 5:1	R > S 8: 0–70 $p < 0.001$		R < S 26: 130–380 $p < 0.001$	R > S 9: 190–270 $p < 0.001$			
PP 6:1	R > S 10: 0–90 $p < 0.001$		R < S 29: 130–410 $p < 0.001$	R > S 28: 140–410 $p < 0.001$	R > S 5: 180–220 $p < 0.001$		
PP 7:1	R > S 9: 0–80 $p < 0.001$			R > S 29: 90–370 $p < 0.001$	R > S 19: 160–340 $p < 0.001$	R > S 7: 130–190 $p < 0.005$	
PP 8:1				R < S 16: 190–340 $p < 0.001$	R > S 34: 90–420 $p < 0.001$	R > S 19: 160–340 $p < 0.001$	
PP 9:1				R < S 29: 120–400 $p < 0.001$	R > S 23: 90–310 $p < 0.001$	R > S 30: 130–420 $p < 0.001$	R > S 11: 160–260 $p < 0.001$

Table 6.1: Results for the comparisons between real (R) and surrogate (S) datasets for the phase-amplitude (PA) and phase-phase (PP) couplings measured with the SFT. Each cell indicates the dataset yielding the highest coupling strength, the number of successive significant windows, the corresponding time interval in ms with respect to the stimulus onset, and the p-value of the corresponding permutation test. Coupling coefficients are shown in the left column for the phase-phase couplings. Blank cells correspond to irrelevant combinations of bands and coupling coefficients.

	1–4 Hz 4–8 Hz	4–8 Hz 8–12 Hz	4–8 Hz 15–25 Hz	4–8 Hz 35–45 Hz	4–8 Hz 45–55 Hz	4–8 Hz 55–65 Hz	4–8 Hz 65–75 Hz
PA	IC > NC 4: 370–400 $p < 0.05$	IC < NC 18: 310–480 $p < 0.001$	IC < NC 8: 310–380 $p < 0.005$	IC < NC 16: 250–400 $p < 0.001$	IC < NC 14: 330–460 $p < 0.005$	IC < NC 20: 290–480 $p < 0.001$	IC < NC 22: 190–400 $p < 0.001$
PP 4:3	IC < NC 14: 350–480 $p < 0.001$	IC > NC 29: 200–480 $p < 0.001$					
PP 3:2	IC < NC 4: 350–380 $p < 0.01$	IC > NC 28: 190–460 $p < 0.001$					
PP 2:1	IC = NC	IC < NC 35: 130–470 $p < 0.001$	IC = NC				
PP 3:1	IC > NC 4: 510–540 $p < 0.001$		IC = NC				
PP 4:1	IC = NC		IC < NC 5: 350–390 $p < 0.001$				
PP 5:1	IC = NC		IC < NC 9: 320–400 $p < 0.001$	IC > NC 9: 300–380 $p < 0.001$			
PP 6:1	IC = NC		IC = NC	IC > NC 21: 240–440 $p < 0.001$	IC > NC 5: 480–520 $p < 0.001$		
PP 7:1	IC = NC			IC = NC	IC > NC 9: 180–260 $p < 0.001$	IC > NC 17: 340–500 $p < 0.001$	
PP 8:1				IC = NC	IC > NC 5: 330–370 $p < 0.005$	IC > NC 18: 150–320 $p < 0.001$	
PP 9:1				IC < NC 18: 230–400 $p < 0.001$	IC = NC	IC > NC 37: 160–520 $p < 0.001$	IC > NC 5: 360–400 $p < 0.001$

Table 6.2: Results for the comparisons between IC and NC conditions for the phase-amplitude (PA) and phase-phase (PP) couplings measured with the SFT. Each cell indicates the condition yielding the highest coupling strength, the number of successive significant windows, the corresponding time interval in ms with respect to the stimulus onset, and the p-value of the corresponding permutation test. Coupling coefficients are shown in the left column for the phase-phase couplings. Blank cells correspond to irrelevant combinations of bands and coupling coefficients.

	1–4 Hz 4–8 Hz	4–8 Hz 8–12 Hz	4–8 Hz 15–25 Hz	4–8 Hz 35–45 Hz	4–8 Hz 45–55 Hz	4–8 Hz 55–65 Hz	4–8 Hz 65–75 Hz
PA	IC > NC 5: 310–350 $p < 0.05$	IC < NC 9: 150–230 $p < 0.001$	IC = NC	IC < NC 15: 230–370 $p < 0.005$	IC = NC	IC = NC	IC = NC
PP 4:3	IC < NC 14: 50–180 $p < 0.001$	IC > NC 14: 50–180 $p < 0.001$					
PP 3:2	IC < NC 4: 150–180 $p < 0.005$	IC > NC 24: 100–330 $p < 0.001$					
PP 2:1	IC = NC	IC < NC 29: 50–330 $p < 0.001$	IC > NC 4: 450–480 $p < 0.001$				
PP 3:1	IC > NC 4: 490–520 $p < 0.001$		IC = NC				
PP 4:1	IC = NC		IC = NC				
PP 5:1	IC = NC		IC < NC 6: 210–260 $p < 0.001$	IC > NC 5: 170–210 $p < 0.001$			
PP 6:1	IC = NC		IC = NC	IC > NC 13: 120–240 $p < 0.001$	IC > NC 6: 190–240 $p < 0.001$		
PP 7:1	IC = NC			IC = NC	IC > NC 20: 50–240 $p < 0.001$	IC > NC 8: 190–260 $p < 0.001$	
PP 8:1				IC < NC 11: 150–250 $p < 0.001$	IC > NC 9: 200–280 $p < 0.001$	IC > NC 7: 170–230 $p < 0.001$	
PP 9:1				IC < NC 11: 160–260 $p < 0.001$	IC = NC	IC > NC 13: 120–240 $p < 0.001$	IC > NC 6: 100–150 $p < 0.005$

Table 6.3: Results for the comparisons between IC and NC conditions for the phase-amplitude (PA) and phase-phase (PP) couplings measured without the SFT. Each cell indicates the condition yielding the highest coupling strength, the number of successive significant windows, the corresponding time interval in ms with respect to the stimulus onset, and the p-value of the corresponding permutation test. Coupling coefficients are shown in the left column for the phase-phase couplings. Blank cells correspond to irrelevant combinations of bands and coupling coefficients.

Identification of the Outcome of Atrial Fibrillation Ablation

7

This chapter presents an application of the harmonic frequency tracker (HFT), the extension of the single frequency tracker (SFT) for harmonic components introduced in Chapter 4. More precisely, the HFT was applied to electrocardiographic (ECG) signals in the context of a study about the outcomes of catheter ablation for atrial fibrillation. In the following sections, atrial fibrillation and catheter ablation are briefly introduced [163], and then all aspects of this study are detailed. This study was presented in two publications [68, 69].

7.1 Introduction

Atrial fibrillation (AF) [164] is the most common cardiac rhythm disorder (arrhythmia). In AF, the upper chambers of the heart, the atria, activate irregularly and do not beat at the normal pace, the so-called sinus rhythm. Clinical guidelines classify AF in three categories depending on its duration [165–167]: paroxysmal, persistent and permanent AFs. The arrhythmia is called paroxysmal when it terminates spontaneously and lasts less than seven days. When sustained for longer than seven days, AF is designated persistent. Frequently, it is not self-terminating and thus may require a clinical intervention (called cardioversion) to restore the sinus rhythm. Long-standing persistent AF, a subcategory of persistent AF, lasts for more than one year and usually leads to permanent AF. In this case, cardioversion has failed to terminate the arrhythmia, or has even not been attempted. AF results in impaired atrial mechanical activity, irregular ventricular activation, fast heart rate, and decreased cardiac output [165, 167]. These dysfunctions can cause several symptoms such as severe palpitations, angina, dyspnea, decreased exercise tolerance, and more generally decreased quality of life [168, 169]. Although not directly life-threatening, AF was associated with substantial morbidity and mortality [170]. Indeed, this tachyarrhythmia has a potential tendency to precipitate and complicate heart failures. Furthermore, due to the impairment of the atrial mechanical activity, AF promotes the formation of blood clots, which in turn increases the risks for stroke and thromboembolic events. As stated previously, AF is the most common arrhythmia in clinical practice. It has been estimated that this condition affects 4.5 million people in the European Union and 2.2 million in the United States [167]. In addition, it has been projected that 15.9 million will be affected in America by 2050, due to various factors such as the AF prevalence increasing with age and the aging of the population [171]. Consequently, as the health expenditures related to this arrhythmia are already important, the financial costs of AF management will become a critical matter in the future [172].

Several guidelines have been proposed to manage this tachyarrhythmia [165–167, 173]. Namely, they provide a framework for treating AF in clinical settings. Possible treatments in-

clude pharmacological and electrical cardioversions, low-energy pacing protocols by means of implanted atrial pacemakers, and surgical and catheter-based ablations. In particular, stepwise radiofrequency catheter ablation (step-CA) has become the treatment of choice for the restoration of sinus rhythm in patients with persistent AF [167]. Due to the evolutive nature of atrial fibrillation [174], electrical and structural remodeling of the atria usually takes place with AF sustainance, providing diffuse pathological substrate for further AF perpetuation. Sites with abnormal high frequency discharges, described as complex fractionated electrograms, and regions showing structural discontinuities in the left atrium (LA) are currently believed to be involved in AF maintenance. Therefore, additional strategies to improve ablation success rate are required. One of them is the creation of additional linear lesions. However, as the amount of ablation needed to obtain long term sinus rhythm is unknown, the success rate of step-CA appears limited. Furthermore, to achieve high success rate, the ablation procedure typically lasts several hours, which increases the risk of complications. In addition, the large electrical scars that are created during the step-CA can play a major role as a substrate for subsequent arrhythmias and probably increase the need for supplemental procedures. There is therefore a strong interest in predicting the outcome of step-CA to decrease both procedural time and ablation extent. To achieve this goal, a better understanding of AF dynamics and organization is clearly needed.

Several invasive and non-invasive techniques have been proposed to characterize the AF dynamics, quantify AF organization and predict the outcome of catheter ablation. It is clear that, ultimately, non-invasive methods present the most benefit as they are easier to implement and less restrictive for the patient. Such methods are typically based on the standard 12-lead surface ECG after ventricular activity cancellation [175]. A simple example is the AF cycle length (AFCL) which is computed by analyzing the main frequency of the ECG fibrillatory waves (F-waves) [176]. It was shown that patients with longer baseline ECG AFCL have more successful outcomes for single LA catheter ablation [177]. In the case of persistent AF, some patients require bi-atrial ablation. For instance, in a recent study [178], whenever the intracardiac AFCL (measured invasively in this case) was shorter in the right atrial appendage (RAA) than in the left atrial appendage (LAA) following LA ablation, right atrium (RA) ablation was performed. This study established that baseline intracardiac AFCL was the best predictor of a successful procedure, but without distinguishing between patients terminating after LA ablation only or after bi-atrial ablation. In two other investigations [179, 180], the F-wave amplitude was shown to predict the procedure outcome. It should be noted that, in the first of these studies, the analysis was performed manually on individual leads. By contrast, in the second one, the proposed method was automatic and took into account the inter-lead variability to increase prediction accuracy. In both investigations, only LA ablation was conducted. Techniques based on AF spectral contents have also been developed to assess AF complexity. Everett et al. [181] introduced the first organization index based on power spectral density estimation to quantify AF organization. In practice, this index is computed as the ratio of the area below the fundamental and harmonic peaks to the total area of the magnitude spectrum. It was used to enhance the success rate of electrical cardioversion by delivering shocks during periods of high organization. In addition, this index was coupled to multivariate frequency analysis to distinguish persistent AF from long-lasting AF successfully [182, 183]. Another approach characterized AF organization based on the phase relations between fundamental and harmonic components [184]. In persistent AF, the fibrillatory waveforms may fluctuate widely over time, and these techniques may not always be capable to track sharp variations [184, 185]. Indeed, to a certain extent, all these methods implicitly assume that the fundamental frequency of AF activity remains constant. Also, at the present time, no study was able to distinguish patients that can be successfully treated by LA ablation only from patients that require bi-atrial ablation. Therefore, new measures of organization based on adaptive frequency tracking are proposed to tackle these issues and provide a global assessment of AF dynamics. Furthermore, these measures can help to predict the outcome of the ablation procedure like the intracardiac AFCL [178], but on the basis of surface ECG recordings only.

This chapter is organized as follows. The materials and methods are described in Section 7.2. The results are presented in Section 7.3, followed by a discussion in Section 7.4. Last, it should be noted that, thereafter, the first harmonic component is considered to be the fundamental component.

7.2 Materials and Methods

7.2.1 Ablation Protocol

The step-CA protocol, illustrated in Figure 7.1, consisted in pulmonary veins isolation (PVI), defragmentation of complex fractionated atrial electrograms (CFAEs), LA linear ablations (roof and mitral isthmus), ablation of RA CFAEs and linear ablation of the cavotricuspid isthmus. The protocol was ended when AF was terminated into sinus rhythm or atrial tachycardia. Non terminated AFs were cardioverted electrically. After restoration of sinus rhythm, verification of conduction block (PVI and lines) was performed, and additional ablations were delivered to achieve a complete block when needed.

7.2.2 Patient Population and Electrophysiological Measurements

The study group consisted of 17 consecutive male patients with persistent AF. Persistent AF was defined as continuous AF lasting longer than four months, resistant to either pharmacological or electrical cardioversion. The patients, depending on the outcome of their respective ablation procedure, were classified into three groups:

Left terminated (LT) Patients that terminated into sinus rhythm or atrial tachycardia during CFAEs or linear ablation in the LA ($N = 11$).

Right terminated (RT) Patients that terminated into sinus rhythm or atrial tachycardia during ablation of the RA ($N = 2$).

Not terminated (NT) Patients that remained in AF and were cardioverted electrically ($N = 4$).

The characteristics of the patient population are reported in Table 7.1.

All patients had effective anticoagulation therapy for more than one month. All antiarrhythmic drugs, with the exception of amiodarone and beta-blockers, were discontinued five half-lives before the ablation procedure, which was performed in general anesthesia. The following catheters were introduced via the right femoral vein: a 3.5-mm cooled-tip ablation catheter for mapping and ablation, a steerable decapolar catheter within the coronary sinus, a circumferential duodecapolar Lasso catheter within the left atrium after transseptal access, and a quadripolar catheter into the RAA for continuous monitoring. The heart electrical activity was also recorded with standard 12-lead surface ECG, except that chest lead V6 was placed in the back (V6b) to improve the recording of the antero-posterior activity of the atria. Surface ECG and endocardial electrograms (EGMs) were continuously monitored and recorded for offline analysis at 2-kHz sampling rate (Axiom Sensis XP, Siemens). For each patient, ECGs and EGMs were separated in 10-s epochs arranged in chronological order. These epochs were not necessarily consecutive as signals recorded during ablation were discarded. On average, 11 ± 3 minutes were recorded at baseline for each patient.

7.2.3 Signal Processing

In order to analyze the atrial activity, ventricular activity was canceled from the 12-lead surface ECG recordings with the single-beat method [175]. This method treats QRS complexes and T-waves separately: the T-wave cancellation is based on a dominant T-wave approach, and the

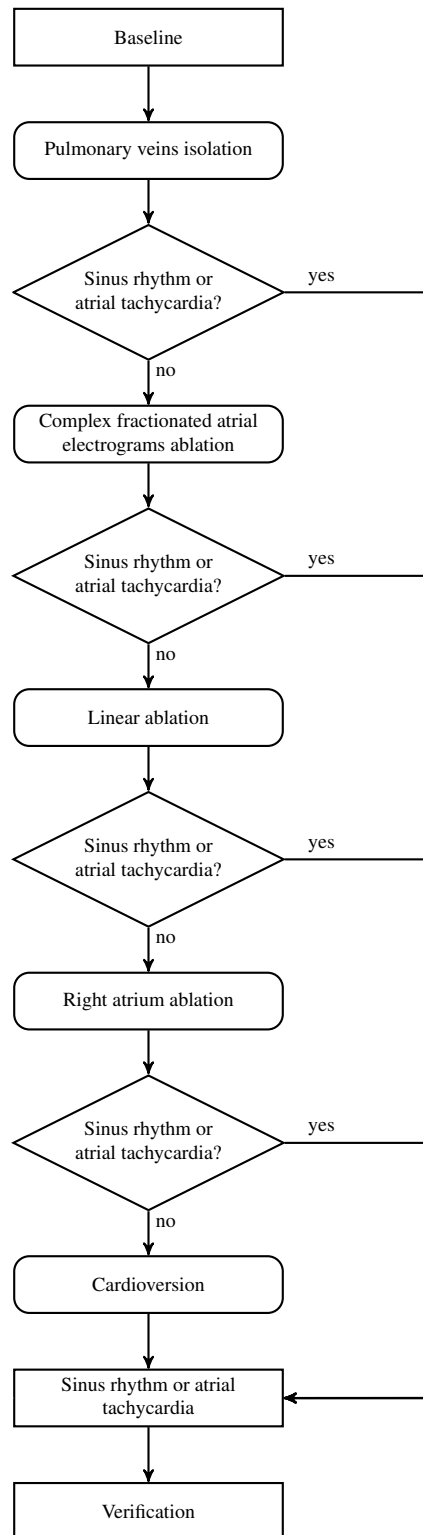


Figure 7.1: Stepwise radiofrequency catheter ablation protocol.

Characteristic	Left terminated <i>N</i> = 11	Right terminated <i>N</i> = 2	Not terminated <i>N</i> = 4
Age (years)	59 ± 6	58 ± 0.7	62 ± 5
Duration of AF (years)	8 ± 6	7 ± 0	4 ± 2
Duration of substained AF (months)	17 ± 8	11 ± 2	39 ± 11
Body mass index (kg/m ²)	32 ± 6	28 ± 4	27 ± 7
Left atrial volume (ml)	179 ± 26	178 ± 87	164 ± 33
Left ventricular ejection fraction (%)	44 ± 10	58 ± 4	54 ± 13
Sites of AF termination			
Left atrium			
Roof	2		
Left atrial appendage	1		
Coronary sinus	2		
Mitral isthmus	6		
Right atrium			
Cavotricuspid isthmus		1	
Right atrial appendage		1	

Table 7.1: Patient characteristics. Means ± standard deviations are indicated when relevant.

QRS complexes are suppressed by estimating the atrial activity with a weighted sum of sinusoids. The single-beat method has two main advantages. First, the length of the ECG recordings is not an issue. Second, it does not introduce any discontinuity and the QRS residues are reduced in the canceled signals. This method was thoroughly tested on both synthetic and real ECG signals and compared favorably to other widely-used techniques. The ECG recordings with canceled ventricular activity were then downsampled at 50 Hz since this investigation focused on the first two harmonic components and previous studies have shown that it is very uncommon for the instantaneous fundamental frequency of the atrial activity to rise above 10 Hz [186].

In order to assess AF organization, the harmonic components of the atrial activity were extracted from the ECG recordings. The approach introduced previously [184] for characterizing F-wave complexity using phase information performs this extraction in two steps. First, the fundamental frequency is determined by locating the largest peak in the interval from 3 to 12 Hz in a power spectral density estimate. Then, the harmonic components are extracted using linear time-invariant band-pass filters centered on the identified spectral peaks. In the study presenting this method, the phases of these components were estimated over non-overlapping adjacent blocks, and the pairwise phase differences were used to characterize AF waveforms. This technique is based on the observation that, for perfectly regular oscillations, the harmonic frequencies are integer multiples of the fundamental one, whereas it is not the case for irregular oscillations. Therefore, the phase differences between fundamental and harmonic components should be close to zero (modulo 2π) for regular (or organized) oscillations. On the contrary, the absolute values of these phase differences should increase for more complex oscillations. The AF organization can thus be quantified with these phase differences. As mentioned previously, one limitation of this approach is that, in practice, the fundamental frequency is not constant. Therefore, there is a tradeoff for setting the bandwidths of the band-pass filters. Indeed, with too narrow bandwidths, the harmonic frequencies may wander outside of the passbands, while too wide bandwidths may let too much noise through and thus disrupt the phase extraction. In the present study, to alleviate this problem caused by the fixed band-pass filters, the HFT was used to estimate the instantaneous fundamental frequency and extract the harmonic components. As this adaptive scheme is presented in Section 4.4 in detail, only its key aspects are repeated here. First, the HFT is

composed of two parts, time-varying band-pass filters and adaptive mechanisms, as illustrated in Figure 4.13. Each filter extracts one of the harmonic component, and the corresponding adaptive mechanism provides an estimate of the fundamental frequency based on this filtered component. At last, a global frequency estimate is computed by combining all estimates with a weighting procedure. Furthermore, this algorithm attempts to maximize the oscillatory behavior at its outputs since it is derived from the SFT, an adaptive scheme based on the complex discrete oscillator equation. The band-pass filters have also two very important properties, namely unit gain and zero phase at the central frequency. Consequently, the extracted harmonic components are not distorted, which is of the utmost importance when investigating their phase relations. One last point is that the HFT is designed to process complex-valued signals. Thus, it was applied to the analytic representations of the ECG recordings, which were computed with the discrete Hilbert transform (DHT) [34]. In this study, the HFT was always set to track the first (or fundamental) and second harmonic components. In addition, the bandwidth and update parameters of this adaptive tracking scheme were set to $\beta = 0.95$ and $\delta = 0.95$.

7.2.4 Organization Measures

The first organization measure considered in this study is the intracardiac AFCL. Local activation times were extracted from the bipolar recordings by detecting their maximum positive peaks using sliding windows of 150 ms. False detections were removed using temporal and amplitude thresholding. The intracardiac AFCL was measured in the LAA, RAA and coronary sinus (CS). The AFCL was also computed for each precordial lead (V1-V6b) as the inverse of the dominant frequency. The dominant frequency was estimated as the frequency of the largest peak between 3 and 15 Hz in the power spectral density estimate of each lead. This estimate was computed with Welch's method [1, 2]. Longer AFCLs typically reflect more organized AF.

The second measure is the organization index (OI) [181]. In this study, it was obtained with the following steps. First, the input signal was band-pass filtered between 1 Hz and 10 Hz. Then, the power spectrum of the resulting filtered signal was computed with the discrete Fourier transform (DFT), and the largest peak was located. Last, the OI was computed as the ratio of the power in a 1-Hz band centered on this peak to the total power in the spectrum.

Similar to the OI, an adaptive OI (AOI) was defined as the ratio of the power of the first two harmonic components extracted with the HFT to the total power of the signal. The local estimates of these powers were computed by low-pass filtering the squared input and outputs with a Hamming window of length 101 [31]. This yielded an organization measure localized in time. Whenever a global value was needed, the local AOI was averaged over the whole signal duration.

The filtered outputs of the adaptive scheme were also used to define another organization measure derived from the phase difference between harmonic components. This measure is closely-related to one proposed in [184] which is based on the observation that the difference of the phases of two oscillations is an indicator of their synchronization. A typical method for extracting the phase information is the Hilbert transform and its associated analytic representation [23]. However, interpreting phase information extracted from broad-band signals can be problematic. Indeed, it has been shown that proper estimation of phase parameters can only be performed on narrow-band signals [27, 30]. Applying the HFT ensures that the phase information of the extracted harmonic components is meaningful. Thus, a suitable evaluation of the complexity of the AF under study can be performed. This organization measure was computed according to the following steps. The first and second harmonic components were extracted with the tracking algorithm, and their phases were computed with the DHT. The phase of the second component was divided by two to ensure that both quantities were comparable. Once the phase difference was computed, its slope was locally estimated by fitting a polynomial of degree one to centered sliding windows of odd length L , which was set to $L = 101$ in this study. The closer

the slope is to one, the more organized is the signal of interest. Then, the phase difference slope variance (PDSV) was computed as the variance of this local estimated slope. Small values of this measure indicate data with low complexity, and thus more organized oscillations.

The AOI and PDSV are illustrated with two examples: one with a synthetic signal and one with real surface ECG data. In the first example, the various steps of the analysis procedure (harmonic frequency tracking and extraction of the organization measures) are also described. The synthetic input signal is composed of two frequency-modulated (FM) harmonic components embedded in additive white noise:

$$x[n] = A_1 \sin\left(2\pi f_0 n + \frac{f_d}{f_m} \cos(2\pi f_m n)\right) + A_2 \sin\left(2\pi 2f_0 n + \frac{2f_d}{f_m} \cos(2\pi f_m n) - \frac{4\pi}{500} n\right) + v[n], \quad (7.1)$$

where $A_1 = 1$ and $A_2 = 0.5$ are the amplitudes, $f_0 = 0.1$ is the fundamental frequency, $f_d = 0.02$ is the frequency deviation, $f_m = 0.002$ is the modulation frequency, and $v[n]$ is an additive white Gaussian noise with a SNR of 10 dB. The second component also includes a linear phase drift. Figure 7.2 shows the input signal, the extracted harmonic components, the estimated fundamental frequency, and the corresponding organization measures. Owing to the HFT, the fundamental frequency and the components are precisely extracted which, in turn, leads to reliable organization measures. In the second example, a short excerpt of a real signal of atrial activity after ventricular cancellation recorded from lead V1 was used. It is illustrated in Figure 7.3 alongside the corresponding AOI and phase difference slope. In the beginning, the signal is poorly organized as indicated by the low AOI and high PDS. However, around the 5-s mark, it becomes more organized as the AOI rises and the PDS decreases.

All the considered organization measures (AFCL, OI, AOI and PDSV) were computed on all available 10-s epochs recorded at baseline for all leads. Then, they were averaged over the epochs of each subject. This resulted in one value for each measure, subject and lead. The intracardiac AFCLs recorded with electrodes located in the LAA, RAA and CS were computed in the same manner. It is also important to mention that the first two patients in LT group were discarded when analyzing lead V6b. Indeed, this lead was recorded at the standard position (V6) for these two patients.

7.2.5 Statistical Analysis

Two different analyses were performed. First, the relations between surface and intracardiac baseline AFCLs were investigated with Pearson's correlation coefficient [76]. For this purpose, the correlation coefficient was computed for each combination of leads (V1-V6b) and intracardiac electrodes (LAA, RAA and CS). Second, the organization measures (AFCL, OI, AOI and PDSV) were compared across the three groups (LT, RT and NT) for all leads within the analysis of variance (ANOVA) framework [187]. In each case where the ANOVA declared a significant effect, follow-up pairwise *t*-tests were performed in order to distinguish which groups were different [188].

7.3 Results

The results of the comparison between mean AFCLs of chest leads V1 to V6b and intracardiac LAA, RAA and CS AFCLs at baseline are shown in Figure 7.4. The correlation between RAA and chest leads was maximal for V1 ($r = 0.96$) and progressively dropped until V5 ($r = 0.26$), with a moderate rise for V6b ($r = 0.62$). LAA AFCL showed the opposite pattern with the highest correlation in V6b ($r = 0.95$) and the lowest one in V2 ($r = 0.26$). The correlation between CS and surface AFCLs was similar to the LAA one, with two maxima at V1 and V6b.

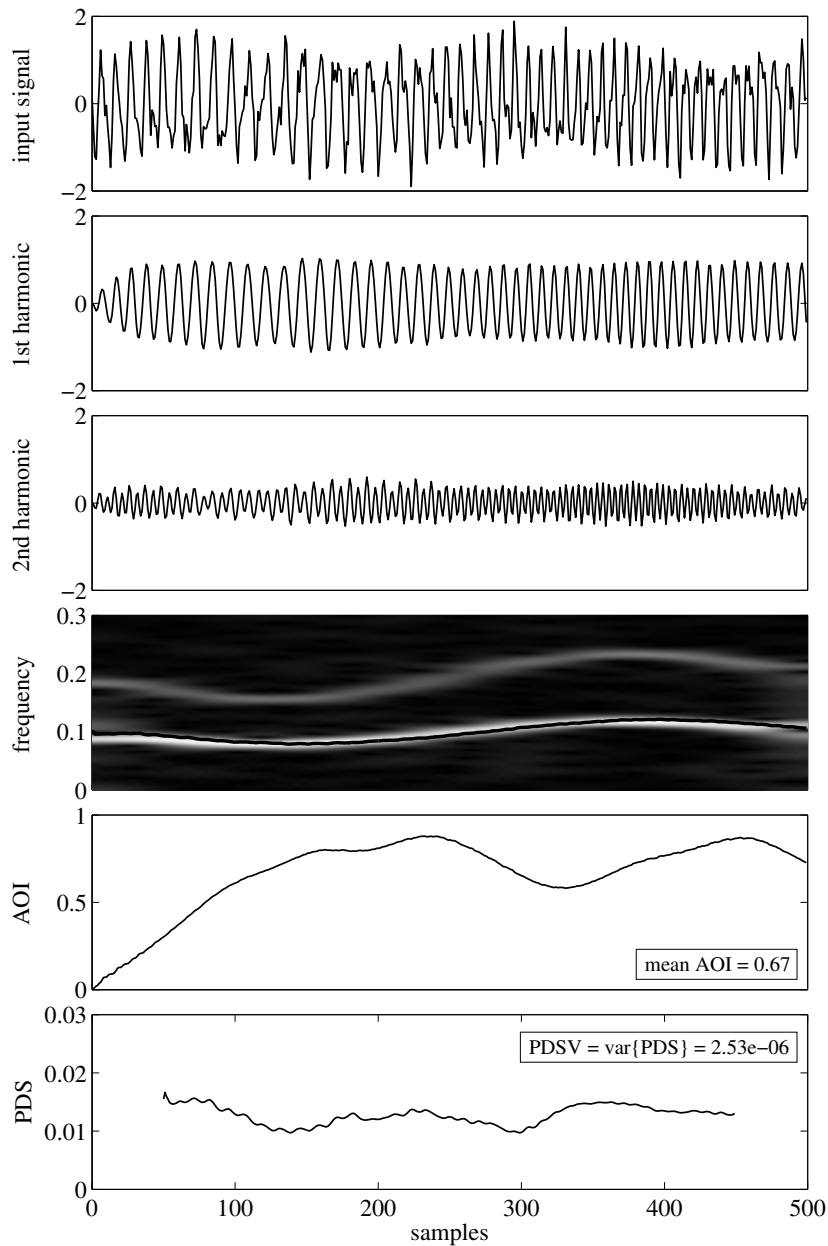


Figure 7.2: Harmonic frequency tracking example. From top to bottom are plotted the input signal, the extracted first and second harmonic components, the estimated fundamental frequency on top of a short-time Fourier transform of the input signal, the AOI and the phase difference slope (PDS). The first and last 50 samples of the PDS were not computed as they required the knowledge of the phase difference for $n < 0$ and $n > 500$.

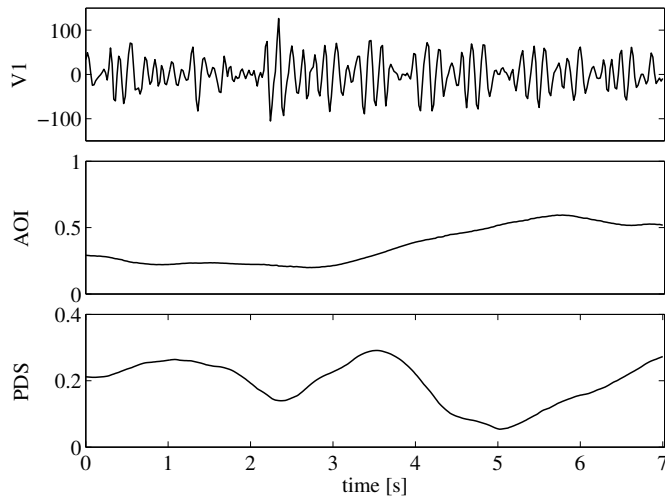


Figure 7.3: Example of real signal of atrial activity. From top to bottom are plotted a signal from lead V1 after ventricular cancellation, the corresponding AOI and the phase difference slope (PDS). Low AOI and high PDS indicate poor organization in the beginning. By contrast, around the 5-s mark, the signal becomes more organized as the AOI increases and the PDS decreases.

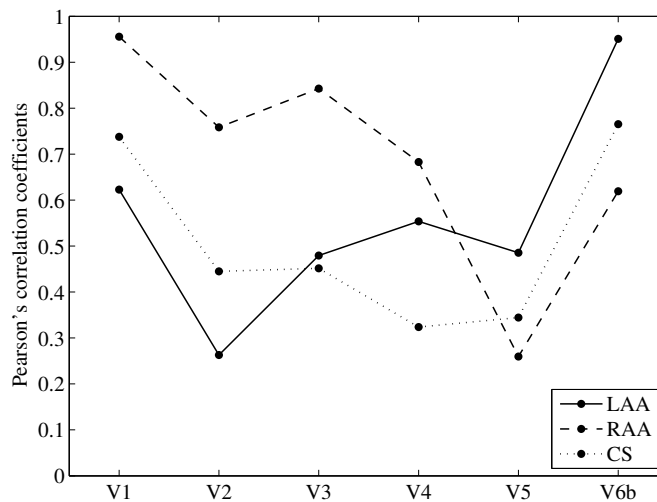


Figure 7.4: Pearson's correlation coefficients between chest leads V1 to V6b and intracardiac LAA, RAA and CS AFCLs at baseline.

Boxplots showing the considered organization measures (AFCL, OI, AOI and PDSV) for the LT, RT and NT groups at baseline are plotted in Figure 7.5 for all leads. Overall, the measures indicated that the AFs that terminated during LA ablation were more organized. However, there were a few exceptions: AFCL for lead V5, OI for leads V2 to V6b, and PDSV for lead V5. Similarly, in all significant ANOVAs, except the one for the OI at lead V4, the AFs terminating after LA ablation were less complex. The means and standard deviations of all organization measures are shown in Figure 7.6 for all leads, alongside follow-up pairwise *t*-tests when relevant. The *t*-tests highlighted two very interesting combinations, namely the AOI and PDSV measured at lead V1 and V3 respectively. In these two cases, the patients in LT group were significantly different from the patients in the two other groups. Therefore, these two measures might prove helpful for predicting the outcome of step-CA *a priori*. Furthermore, in terms of PDSV, the RT group was significantly different from both the LT and NT groups at leads V4 and V5. A counter-intuitive case involved the OI and lead V4 where LT group was significantly less organized than the other two. This result might be due to the limited number of patients included in this study. In all the other cases where the ANOVA revealed significant differences, the *t*-tests could only distinguish between two groups. In practical situations, this is less useful for predicting ablation outcome.

7.4 Discussion

The correlation of the ECG AFCLs and the multi-site intracardiac AFCLs has shown that both atria dynamics are reflected in the surface ECG, as reported in previous studies. Holm et al. [186] have shown that the ECG devoid of ventricular activity reflects the dynamics of the intracardiac AFCL. More specifically, their results suggest that RA is the major contribution to the F-waves in lead V1. Matsuo et al. [177] obtained similar results. By contrast, few studies have investigated how the LA fibrillatory dynamics are reflected in the surface ECG. Platonov et al. [189] have recorded simultaneously a standard 12-lead ECG with electrograms from the RAA and LAA and have concluded that the LA dynamics is also present in the surface ECG, including V1. The study from Petrutiu et al. [190] was extended by using posterior leads. Their results confirmed that lead V1 reflects mostly the RA dynamics, and established a strong correlation between the posterior leads and LA activity. In this study, the CS was also included in addition to LAA and RAA. The analysis of Pearson's correlation coefficients confirmed that the RAA activity is mostly reflected by lead V1 and that the contribution of the LAA can be observed by adding a dorsal electrode. The results for the CS were similar to the ones for the LAA.

Surface and intracardiac AFCL have been used to predict AF termination by step-CA. Haïssaguerre et al. [191] have shown that, during catheter ablation, the prolongation of intracardiac AFCL predicted the conversion of AF into organized tachycardia. Rostock et al. [178] established that baseline intracardiac AFCL was the strongest predictor of ablation success. Interestingly, AFCL before AF termination was predictive on its own, but did not add predictive information once baseline AFCL was selected. Whereas, Matsuo et al. [177] have shown that baseline ECG AFCL was the only predictors of AF termination: patients with longer ECG AFCLs had more successful catheter ablation.

As mentioned previously, current measures of organization (such as OI [181] or the method based on phase relations between harmonics [184]) assume, to a certain extent, that the dominant (or fundamental) frequency of fibrillatory activity remains constant. In persistent AF, the F-waves can fluctuate widely over a short duration, which may affect the reliability of these measures. A solution to this problem is to use the adaptive filters that can take into account such frequency changes. Therefore, the HFT was applied to extract harmonic components. These components were used to compute two organization measures, the AOI and the PDSV. The former is an extension of the classic OI to use the outputs of the HFT, and the latter quantifies organization with the phase difference between harmonic components. These two measures highlighted that,



Figure 7.5: Boxplots of the four organization measures (AFCL in ms, OI, AOI and PDSV) at baseline with respect to the three groups (LT, RT and NT) for all considered electrodes. In each plot, the p-value of the corresponding ANOVA is indicated when significant.

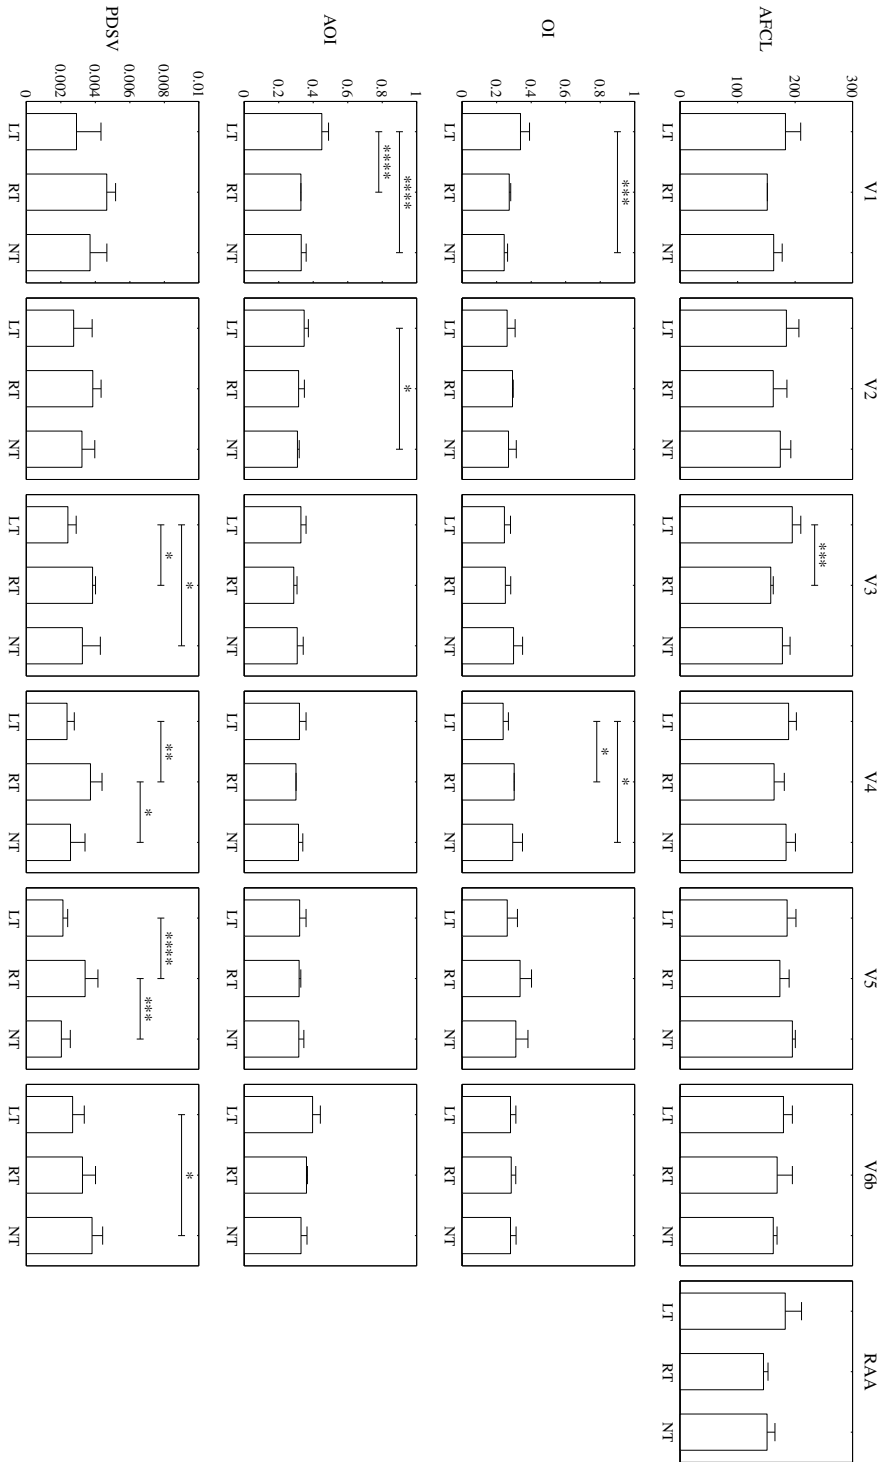


Figure 7.6: Means and standard deviations of the four organization measures (AFCL in ms, OI, AOI and PDSV) at baseline with respect to the three groups (LT, RT and NT) for all considered electrodes. If the corresponding ANOVA declared a significant effect, the significant pairwise comparisons are indicated (*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.005$ and ****: $p < 0.001$).

at baseline, patients in LT group were significantly more organized than the ones in RT and NT groups. Since both approaches are based on the surface ECG, they may pave the way to new techniques for predicting the success rate of step-CA. Such techniques would facilitate the choice of the appropriate treatment for both the physician and the patient. Hence, avoiding ablation procedures with low probability of success would increase the patient quality of life as well as reduce the associated financial costs. Also, the patients in the RT group had significantly different PDSV with respect to patients in the two other groups. This might prove helpful for choosing the most suitable treatment option. Despite these promising results, this study suffers from two main limitations. First, the number of patients is too limited. Second, they are very badly distributed among the three groups. These two limitations might have given rise to the specious results of the OI at lead V4, where patients in group LT had significantly more complex AF than the ones in the other groups. Hopefully, by including more patients in this study, this problem will be solved, without losing the predictive value of the AOI and PDSV. Furthermore, with additional patients and possibly a more balanced distribution among the three groups, the two organization measures based on the HFT may be helpful in the development of a robust classifier for deciding if a patient should undergo a catheter ablation procedure or not. At last, it should also be noted that, although the tracking performance of the HFT are trustworthy, it may need some time to adjust to very abrupt changes in AF dynamics.

Potential Applications of Adaptive Frequency Tracking

8

This chapter briefly presents three potential applications of the single frequency tracker (SFT) and its extensions, the multiple frequency tracker (MFT), the multivariate frequency tracker (MVFT) and the harmonic frequency tracker (HFT), introduced in Chapters 3 and 4. In the first example of application described in Section 8.1, the SFT and HFT were used to estimate the heart rate from a photoplethysmographic signal with poor SNR. In the two other examples, the tracking schemes were applied to retrieve indirectly the instantaneous frequency of components present in related signals. More specifically, the MVFT could recover the pedaling frequency during sprint interval training from a sensor monitoring the concentration of different types of hemoglobin in muscles as shown in Section 8.2, whereas, in Section 8.3, the SFT was used to estimate the respiration frequency from R-R intervals throughout an incremental exercise test for assessing the maximum oxygen uptake. At last, the potential applications of adaptive frequency tracking algorithms are concisely summarized in Section 8.4.

8.1 Heart Rate Estimation from Photoplethysmogram

Photoplethysmography (PPG) is an optical technique to obtain the volumetric measurement of an organ. Its main application is to determine indirectly the heart rate (number of heart beats per unit of time) of a patient by monitoring the oxygen concentration in the blood. This is achieved by measuring the light absorption characteristic of a tissue. Typically, PPG sensors for measuring the heart rate are placed at the finger tip or on the forehead.

A collaboration between the Applied Signal Processing Group (<http://aspg.epfl.ch/>) and Tabrasco Ltd. (<http://www.tabrasco.com/>) aimed at developing an algorithm for estimating the heart rate of archers from such a PPG sensor. This device was designed with two main purposes in mind. First, it could prove useful during training sessions for improving the performance of archers. Second, it could help to make archery competitions more attractive for the audience. Spectators may be more interested when watching television broadcasts of archery contests if they could evaluate the calmness or the stress of an archer with a real-time indication of his (or her) heart rate. There is one important restriction concerning the placement of the PPG sensor: it cannot be located on the upper part of the body. Indeed, it has to be placed below the waist in order to limit as much as possible the inconvenience for the archers. The best location to record the heart rate with a PPG device from the legs is on the posterolateral part of the knee. However, the quality of signals recorded from this location is much lower than from the typical ones (finger tip and forehead). This required more elaborate signal processing techniques than the ones usually implemented for such data. Another restriction was that the complete process-

ing algorithm has to be applicable in real-time, with a delay as small as possible. Therefore, the SFT and HFT were tested to retrieve the heart rate because they are based on order-one filters. In addition, they were shown to perform well in noisy environments (see Chapters 3 and 4). The harmonic extension was considered as a small harmonic component was observed in time-frequency representations.

Here, an example of the signals recorded during the testing phase is presented. In this phase, PPG data were recorded from the knee of healthy subjects who were not practicing archery. Instead, they were instructed to sit on a chair for 60 s, stand for 60 s, walk for 90 s and finally stand for 30 s. This protocol was used in order to assess the performance of the complete procedure for heart rate estimation in different situations. One signal was recorded from each subject and sampled at 30 Hz. In parallel, a Polar[®] sensor was used to provide a reference measure of the heart rate directly from the chest. The SFT and HFT were not applied immediately to the raw PPG signal as it contained a large drift and several artifacts which needed to be suppressed or reduced first. Therefore, some pre-processing was performed. Due to the real-time constraint, the delay introduced by the pre-processing steps was required to be as limited as possible. The first step was to eliminate the drift with a least mean squares (LMS) adaptive filter [192]. Then, the resulting signal was band-pass filtered from 0.75 Hz to 2.5 Hz with an order-30 FIR filter. After the filtering operation, the signal was downsampled at 7.5 Hz in order to recenter the frequency of interest with respect to the frequency range for the adaptive tracking schemes. The band-pass filter ensured that no aliasing was introduced. Figure 8.1 shows a PPG signal before and after applying the pre-processing. Then, the SFT and its harmonic extension were applied with the same set of parameters: $\beta = 0.97$ and $\delta = 0.95$. In addition, the HFT was set to track the first (or fundamental) and second harmonic components. Typically, both adaptive schemes are applied to the analytic representations of real-valued signals. However, it was not possible in this case since the procedure needed to be applicable in real-time. Indeed, the discrete Hilbert transform (DHT), through which the analytic representation is obtained, is computed either via the frequency domain (2.15) (this requires complete knowledge of the signal of interest) or via the Hilbert filter (2.12) which is an IIR filter (truncation of the impulse response would introduce too much delay for acceptable accuracy). Consequently, the SFT and HFT were applied directly to the real signal, and thus the extracted oscillatory components were complex-valued. Nevertheless, real-valued oscillations could be easily obtained by taking the real part. The heart rate estimates provided by the two tracking algorithms from the signal illustrated in Figure 8.1 are shown in beats per minute in Figure 8.2 alongside the estimate from the chest sensor. The three estimated heart rates were in close agreement, except in two occasions. Just after the subject stood up, the SFT and HFT needed more than 30 s compared to the chest estimate to adapt to the heart rate increase. The extremely small signal amplitude following the stance shift hampered the adaption process of both algorithms. A modification of the contact surface between the PPG sensor and the skin is likely to have caused this delay. During the walking phase, the chest sensor provided an erratic and very unreliable heart rate estimate. Indeed, it is highly dubious that such a short walk would increase as much the cardiac activity. The chest sensor seemed to fail completely in this case. Similar outputs were also observed for other subjects. By contrast, the SFT and HFT estimates were much more plausible during the same time interval as they both provided overall trustworthy heart rate estimates. Also, it is worth mentioning that the SFT seemed to react slightly faster to frequency changes than its harmonic extension. This was likely caused by the low power of the second harmonic component.

8.2 Frequency Tracking and Near-infrared Spectroscopy

Near-infrared spectroscopy (NIRS) is a technique for measuring the absorption characteristics of various materials, such as gases, liquids or even biological tissues with respect to the near-

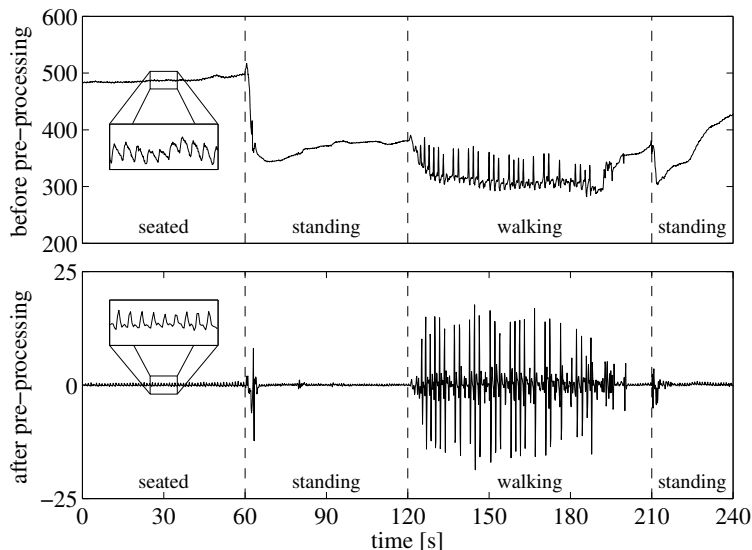


Figure 8.1: PPG signal recorded from the knee. The top graph shows the original signal and the bottom one shows the outcome of the pre-processing (drift cancellation, band-pass filtering and downsampling). Vertical dashed lines indicate the transition between stances. The low-power oscillation corresponding to the heart rate is illustrated in the two insets.

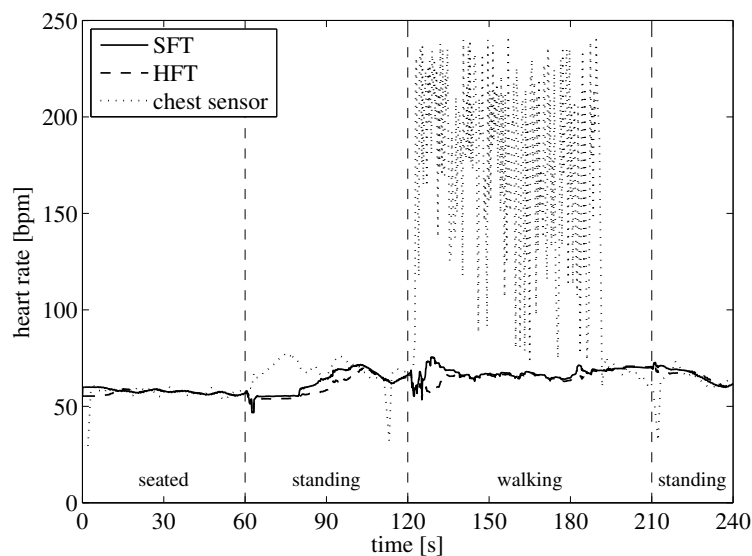


Figure 8.2: Heart rate estimated from a PPG signal with the SFT and HFT as well as from a chest sensor in beats per minute (bpm). Vertical dashed lines indicate the transition between stances.

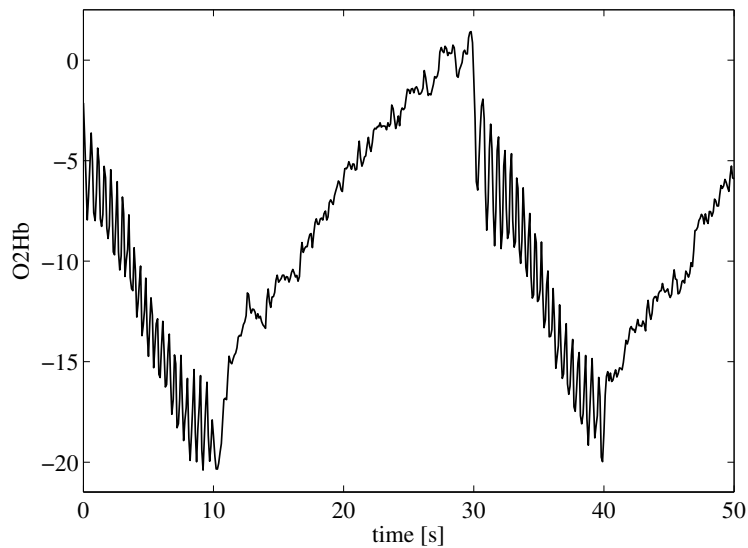


Figure 8.3: Excerpt of an O2Hb signal during sprint interval training in arbitrary units. The large increases and decreases correspond to sprint and recovery intervals respectively.

infrared region of the electromagnetic spectrum. In substance, it can detect specific chemical components. NIRS has been applied to a wide range of fields, such as astronomy, agriculture and medicine. In this section, an application of NIRS to sports science is presented. In more detail, this technique was used to monitor the concentration changes of different types of hemoglobin in muscles during sprint interval training [193, 194].

In the considered study [195], the sprint interval training was performed on a cycle ergometer. After warming up, healthy subjects were asked to alternate between 10-s sprint and 20-s recovery intervals until exhaustion. During the procedure, four different signals were recorded with the NIRS apparatus and sampled at 50 Hz: changes in oxyhemoglobin (O2Hb), deoxyhemoglobin (HHb) and total hemoglobin (tHb), as well as the difference between O2Hb and HHb (dHb). However, it was not possible to record the pedaling frequency directly from the ergometer with the setup used for the acquisition. Nonetheless, the NIRS signals contained, in addition to the strong low-frequency components corresponding to the repetitions of sprint and recovery intervals, the pedaling oscillations in the higher frequencies. Thus, the goal was to perform adaptive frequency tracking in order to retrieve the pedaling frequency. Furthermore, as four signals were available, the MVFT was applied to achieve higher estimation robustness. A short 10-s excerpt of an O2Hb signal recorded during sprint interval training is shown in Figure 8.3. One can observe the slow variations with large amplitude caused by the alternation between sprint and recovery intervals as well as the fast oscillations due to pedaling.

The pedaling frequency could not be estimated straightforwardly as there was too much power in the very low frequency range, and thus it was impractical to apply the adaptive scheme directly. In addition, the starting points of the sprint and recovery intervals needed to be identified. Therefore, some processing was required before estimating the frequency of interest. First, O2Hb, HHb, tHb and dHb signals were low-pass filtered with a cut-off at 0.25 Hz in both forward and reverse directions in order to achieve zero phase distortion [31]. Then, these filtered signals were resampled at 10 Hz and used to detect the edges of the intervals. Afterwards, the very low frequency content was discarded by subtracting the low-pass filtered signals to the original ones. These difference signals were then resampled at 10 Hz as well. Finally, the pedaling frequency

was estimated by applying the MVFT to the resampled differences with the bandwidth and update parameters set to $\beta = 0.9$ and $\delta = \mu = 0.85$. To summarize, the following processing steps were applied to the NIRS data:

1. Low-pass filtering with cut-off at 0.25 Hz.
2. Resampling of the filtered signals at 10 Hz.
3. Detection of sprint and recovery periods on the resampled filtered signals.
4. Differences between original and filtered signals.
5. Resampling of the difference signals at 10 Hz.
6. Adaptive frequency tracking with the MVFT on the resampled differences.

All these processing steps are illustrated in Figure 8.4 for each considered signals (O2Hb, HHb, tHb and dHb). On all graphs, the starting points of sprint and recovery intervals are indicated by green triangles and red squares respectively. There is one more triangle as the subjects were instructed to repeat the sprint-recovery sequence until exhaustion. In the first column, the original signals recorded by the NIRS apparatus are plotted alongside the low-pass filtered ones. The corresponding difference signals obtained by subtracting the low-pass filtered ones to the original ones are plotted in the second column. The third column shows the instantaneous pedaling frequency estimated by the MVFT on top of short-time Fourier transforms of the difference signals. The same frequency trajectory is plotted in the four graphs as the purpose of the MVFT is to estimate the instantaneous frequency of a single oscillatory component present in several signals. In the last column, the adaptive coefficients used by the MVFT to weight the contribution of each signal are shown. It should be noted that all signals were resampled at 10 Hz for illustrative purpose. The multivariate adaptive algorithm highlighted a few interesting points. First, the pedaling began to decrease before the end of the sprint intervals as the subject grew tired. Second, also due to the tiredness, the maximum achieved pedaling frequency decreased with each repetition. Third, the adaptive weights show that, in the beginning, the HHb was the most important signals in terms of frequency estimation, while they were more balanced later. And last but not least, the pedaling frequency, which was not recorded directly, could be recovered.

8.3 Indirect Estimation of the Respiration Frequency

In several practical situations, the frequency of a process that may provide relevant information is not available directly. For instance, it may be too costly, difficult, or even impossible to measure this frequency. Nevertheless, related signals may contain the underlying oscillation, from which the frequency estimation could be indirectly estimated. For example, in the application of adaptive tracking presented in Section 8.2, the pedaling frequency could not be recorded with the acquisition setup. However, it could be retrieved from NIRS signals. Here, the aim was to estimate the respiration frequency during an incremental exercise test for assessing the maximum oxygen uptake ($\text{VO}_2 \text{ max}$) [196, 197] from an R-R interval signal. As mentioned in Section 5.4, R-R intervals are typically used to monitor the heart rate variability, which has two primary causes: the baroreflex and the respiration. Their influences take the form of oscillatory components around 0.1 Hz for the baroreflex and 0.25 Hz for the respiration at rest. During incremental exercise tests however, the respiration rate increases progressively and R-R interval signals are non-stationary. Therefore, classical techniques, such as power spectral density estimation, which require stationarity, perform poorly. This issue can be tackled with time-varying models [198, 199]. But, here, the aim was to estimate the instantaneous frequency of the respiration with the SFT.

The incremental exercise test was performed on a cycle ergometer with the following protocol: 3 min of unloaded pedaling (0 W) at 85 rpm, followed by a load increase of 30 W/min

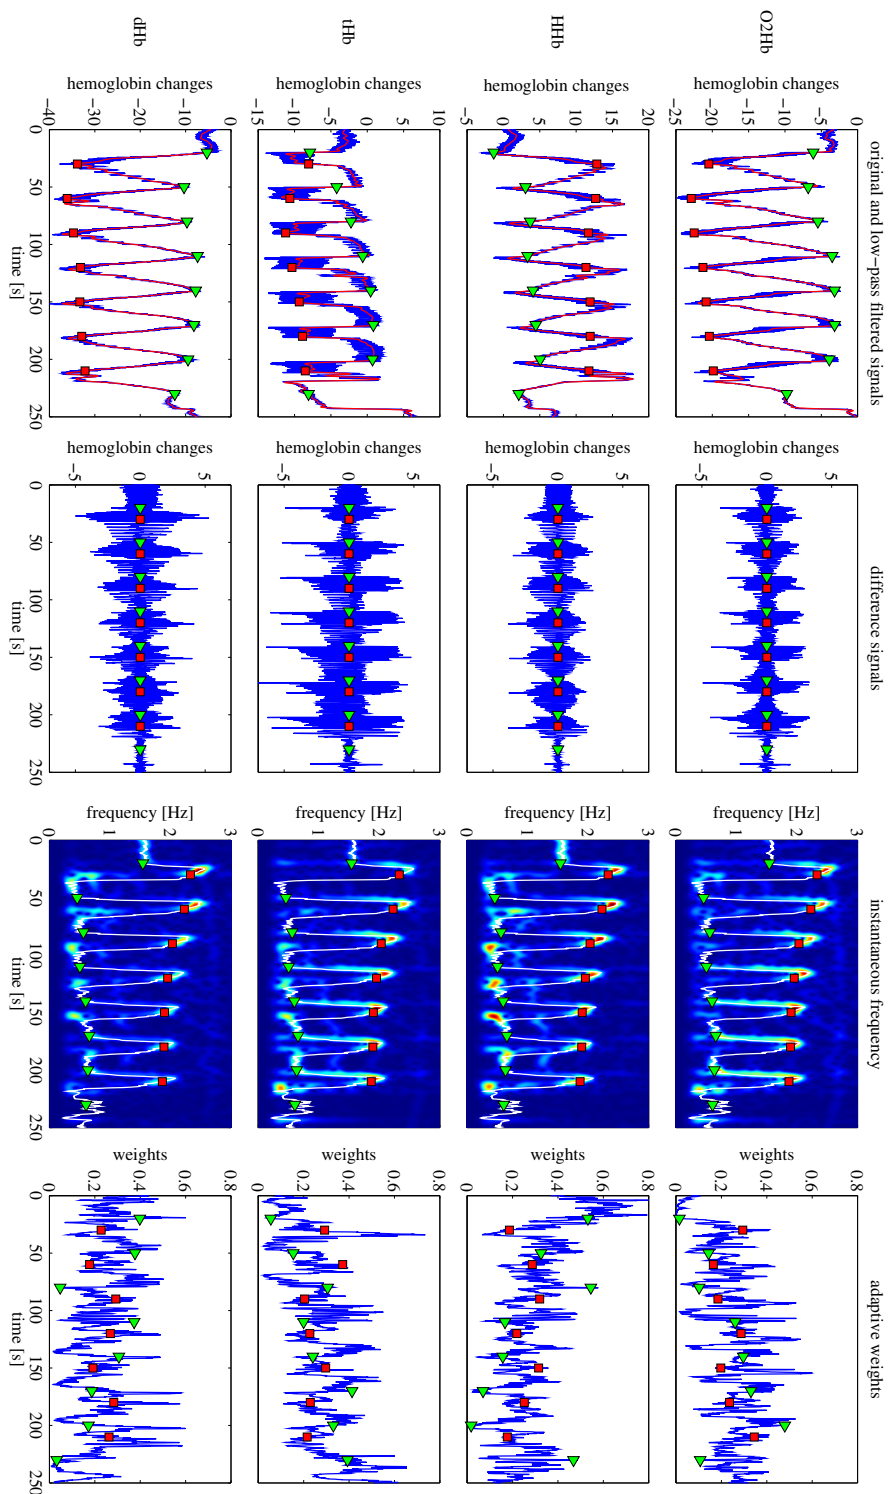


Figure 8.4: Estimation of the pedaling frequency from NIRS data. From left to right are plotted the NIRS (blue) and low-pass filters (red) signals in arbitrary units, the difference signals, the difference signals, the pedaling frequency estimate on top of short-time Fourier transforms of the difference signals, and the adaptive weights used by the MVFT. The green triangles (\blacktriangledown) and red squares (\blacksquare) indicate the starting points of the sprint and recovery intervals respectively.

(1 W increase every 2 s) until exhaustion. Both the R-R intervals and the respiration rate were recorded. Clearly, frequency tracking was not necessary in this case as the respiration frequency was directly recorded. However, it should be considered as a feasibility test since the estimate provided by the SFT could be compared to a reference. Before applying the adaptive tracking, a few pre-processing steps were needed. First, the R-R intervals were regularly resampled at 3 Hz. Then, the drift was suppressed and the resulting signal was high-pass filtered with a cut-off at 0.3 Hz in order to extract the frequency range corresponding to the respiration component. The signal after pre-processing is shown in the top plot of Figure 8.5. Then, the SFT was applied to estimate the respiration frequency with bandwidth and update parameters set to $\beta = 0.95$ and $\delta = 0.9$. The frequency estimate is plotted in Figure 8.6 alongside the reference. The corresponding extracted component is plotted in the bottom graph of Figure 8.5. The SFT estimate and reference match closely, except in the beginning where the initialization of the adaptive algorithm took place and at the largest peak in the end of the recording where the increase was too sharp for the adaptive update. Nevertheless, owing to the SFT, it was still possible to recover indirectly a good estimate of the respiration frequency.

8.4 Summary

Although the three practical applications of adaptive frequency tracking presented in this chapter were taken from the field of sport sciences, the SFT and its extensions are not restricted to this specific field or, more generally, to biomedical signals. Indeed, these examples should be considered as illustrations of the benefits provided by adaptive tracking schemes and of their potential uses. The SFT can be applied whenever an accurate estimate of the instantaneous frequency of an oscillatory component is needed. There is no restriction on the type and source of the considered data. And, since its convergence speed compared favorably with other frequency tracking methods (Section 3.5), this adaptive algorithm is a relevant choice for frequency estimation. Furthermore, it was extended to track simultaneously several oscillations (MFT), and to improve estimation performance either by using the information present in more than one signal (MVFT) or by taking into account the harmonic structure if any (HFT). One of the key reasons to develop the SFT was to tackle the frequency estimation problem in non-stationary data as classical approaches are typically based on a stationarity assumption. This aspect was highlighted in the three presented applications. In particular, the multivariate extension could follow closely the pedaling component in NIRS signals whose instantaneous frequency varied widely.

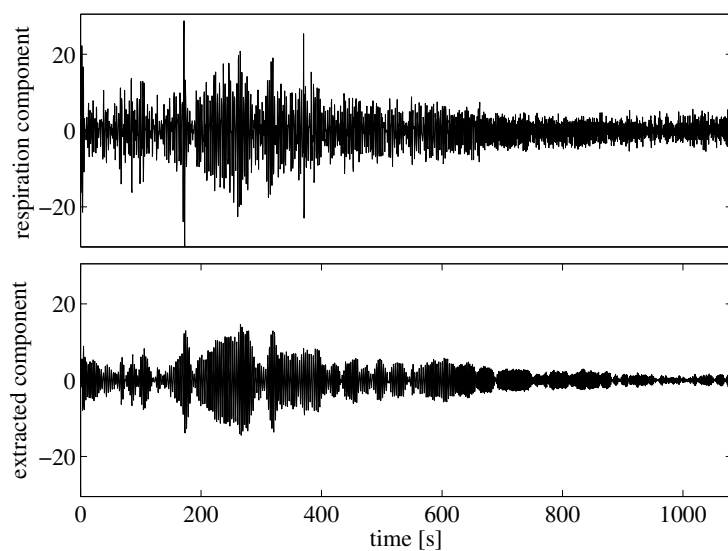


Figure 8.5: Respiration component extracted from R-R intervals. The input and output signals of the SFT are plotted in arbitrary units in the top and bottom graphs respectively.

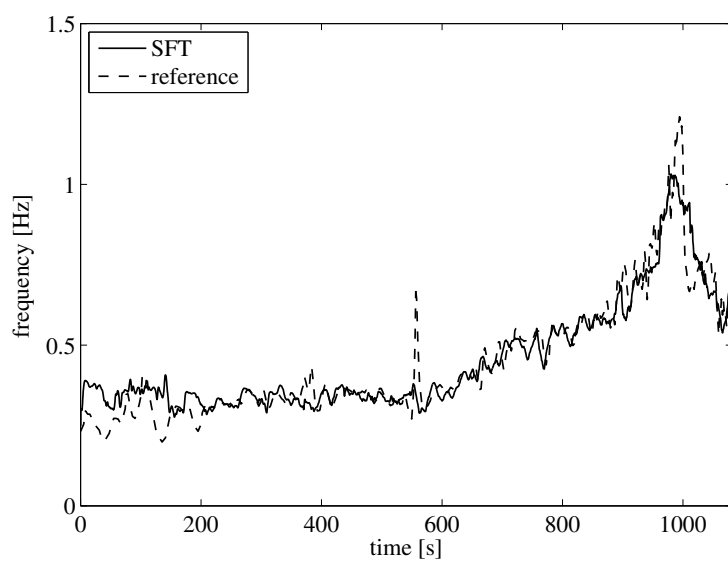


Figure 8.6: Respiration frequency during an incremental exercise test. Both the reference and the estimate provided by the SFT are shown.

9

Conclusion

Frequency is a key aspect in signal processing as its knowledge is of crucial importance in a large number of fields. Consequently, numerous techniques have been proposed over the years in order to retrieve the frequency information from all kinds of signals. A considerable fraction of these approaches focus on oscillatory components whose instantaneous frequency varies over time. Indeed, the temporal evolution of the frequency of an oscillation can provide relevant information about the underlying process. Since instantaneous frequency is such a key concept in signal processing and related fields, the aim of this thesis was to develop efficient schemes for adaptive frequency tracking.

In the first part of the present dissertation, several methods designed to estimate the instantaneous frequency were reviewed. The four methods of this selection based on adaptive filters extracted, in addition to the estimated frequency, the corresponding oscillatory component. This characteristic can be decisive in certain cases. Indeed, the extracted oscillation can provide valuable information or be used in further processing. For instance, this can lead to a proper estimation of the phase, which is only physically meaningful for narrow-band signals. These considerations about the importance of instantaneous frequency estimation and oscillation extraction were the cause of the development of the single frequency tracker (SFT). This adaptive algorithm was designed to track an oscillation buried in noise and estimate its instantaneous frequency. The most distinctive feature of the SFT is that it attempts to maximize the oscillatory behavior at the output. Furthermore, the adaptive band-pass filter used for component extraction has two key properties: unit gain and zero phase at the central frequency. Thus, no distortion was introduced, which is of the utmost importance in some applications, such as phase information estimation. Numerical simulations have shown that the SFT compares favorably to existing techniques for frequency tracking in terms of estimation bias and variance as well as convergence speed.

In many cases of practical interest, the assumption that the signal under study contains a single oscillation embedded in noise is overly optimistic or too restrictive. As a matter of fact, it is common to find several frequency components in a signal. The SFT can only track a single frequency and thus is marginally applicable to such data. An extension of the basic scheme, the multiple frequency tracker (MFT), was developed to handle this issue. It is based on the simple idea of tracking each oscillation with one SFT. However, to ensure that the same component is not followed several times, adaptive all-zero filters were added in order to suppress the crosstalk between trackers. This led to a scheme capable of extracting several oscillations and estimating their instantaneous frequency simultaneously. A related point is raised when the considered signal contains multiple components with harmonic frequencies. The underlying structure can be used to increase robustness and estimation performance since all the frequencies are integer multiples of the fundamental frequency. An extension of the SFT, the harmonic frequency tracker

(HFT), was developed as well in order to take into account this aspect. Numerical simulations confirmed that this extension achieved lower bias and variance than the original algorithm in the presence of harmonic components. Another extension for the SFT was developed based on the observation that in certain fields a phenomenon is frequently investigated with more than one sensor, such as in electroencephalography (EEG) or electrocardiography (ECG). And thus, the oscillation of interest is present in several signals. The multivariate frequency tracker (MVFT) was designed to take into account the additional information provided by all available signals in order to improve the tracking robustness and the overall estimation performance. Multivariate tracking is implemented by estimating the instantaneous frequency for each signal and then by combining all the estimates into a global frequency estimate with a weighting procedure. Owing to the MVFT, it is possible to track an oscillatory component with low amplitude common to several signals, even in cases where single frequency tracking would fail. The gain in estimation variance of the multivariate extension with respect to the SFT was clearly highlighted with numerical simulations.

The SFT and its extensions are causal methods. They only used the knowledge of past samples to estimate the instantaneous frequency. In situations where the data analysis take place offline, the future information can be used in addition to past information in order to achieve better estimation performance. The non-causal frequency estimator (NFE), developed in the course of this thesis, builds upon this observation to enhance an initial frequency trajectory. There is no restriction on the approach used to obtain this initial estimate. For instance, an adaptive frequency tracking scheme or a classical method for power spectral density (PSD) estimation can be used. In fact, the initial estimate can even result from an informed guess. The initial frequency trajectory is refined through an iterative procedure aimed at minimizing a local cost function derived from the complex discrete oscillator equation, similar to the SFT. And the closer the initial estimate is to the true frequency, the faster the algorithm converges. Nevertheless, it should be mentioned that the initialization has to make sense, otherwise the NFE may provide an erroneous estimate. At the cost of real-time application, this method yields an estimated instantaneous frequency without any delay. Furthermore, it also extracts the corresponding component with a non-causal time-varying band-pass filter which does not introduce any delay. These desirable properties however comes at the cost of increased computational complexity and, in large-scale investigations, this load can become prohibitive.

The approaches for instantaneous frequency estimation presented in this dissertation are not ultimate replacements for classical techniques such as methods for PSD estimation or time-frequency analyses. Indeed, these classical approaches are widely spread and used for a reason: they provide relevant information about the frequency content of a signal in numerous situations. Instead, the SFT, its extensions and the NFE should be applied in cases where classical techniques exhibit poor performance or to investigate a problem from a slightly different perspective. In fact, all approaches for instantaneous frequency estimation can be considered as complementary. In certain practical cases, several methods are needed to be applied in conjunction, as the knowledge provided by a single one may not be sufficient.

In the second part of this work, the adaptive frequency tracking schemes developed during this thesis were applied to different types of biomedical data. First, the SFT was used to extract neuronal oscillations and their instantaneous frequency from EEG signals recorded during a visual potential experiment with illusory contour stimuli. The main goal of this study was to investigate the temporal evolution of cross-frequency couplings during visual perception. In particular, the SFT yielded narrow-band oscillations from which proper phase information could be extracted. This information was then used to measure the strengths of phase-amplitude and phase-phase couplings between several frequency bands. The two key outcomes were that the instantaneous frequency in certain bands changed depending on the stimulus type and that the cross-frequency couplings exhibited a complex pattern of increases and decreases depending on the considered band. Furthermore, it should be noted that the SFT revealed the frequency

changes and highlighted the coupling mechanisms more clearly than a traditional filter-bank approach. In another study with ECG data, the HFT was used to develop two organization measures to help predicting the outcome of stepwise catheter ablation of atrial fibrillation (AF). Promising results were obtained. However, the size of the patient population was limited, and thus these results need to be confirmed for additional patients. The practical usefulness of the SFT and its extensions were also illustrated with three examples of signals recorded in the context of sport sciences. Taken together, all these applications highlighted the key role played by the pre-processing steps before frequency tracking. When working with real-world data, one is barely able to apply this kind of approaches directly on the raw signals. Another important point to consider is that, although all the applications presented in this dissertation were based on biomedical signals, the adaptive frequency tracking schemes introduced in Part I are not restricted to this specific class of data. Indeed, frequency is a key concept in a multitude of fields and the developed algorithms only assume that the signals of interest contain oscillatory components. The focus on biomedical signals of this thesis originated primarily from the strong orientation of the Applied Signal Processing Group towards this class of signals.

9.1 Summary of Achievements

The major achievements of this thesis are summarized below.

An Efficient Scheme for Adaptive Frequency Tracking

An adaptive frequency tracking algorithm, the SFT, was developed during this thesis. It can extract an oscillatory component buried in noise and estimate its instantaneous frequency. The most distinguishing feature of the SFT is that it is designed to minimize a cost function derived from the complex discrete oscillator equation, and therefore the oscillatory behavior at the output is maximized. Another relevant aspect of this scheme is its adaptive band-pass filter that has unit gain and zero phase at the central frequency. Thus, no distortion is introduced in the extracted oscillation. Its performance in terms of cost function, bias and variance was thoroughly analyzed with Monte Carlo simulations and, when possible, with theoretical calculations. Furthermore, numerical simulations also showed that the SFT compares favorably to existing methods for instantaneous frequency estimation.

Multiple Frequency, Multivariate and Harmonic Extensions

Three extensions of the SFT were developed in order to tackle three cases of practical interest. First, the MFT was designed to handle signals with more than one oscillatory component. It is capable of tracking several frequencies simultaneously and extracting the corresponding oscillations. Adaptive all-zero filters were added in this extension in order to reduce the crosstalk between the time-varying band-pass filters. Second, sometimes several recordings of the same phenomenon are available. The MVFT was developed in order to achieve better overall performance by taking advantage of this additional information. This extension estimates the frequency of the common component for each signal and then combines all these values into a global estimate through a weighting procedure. The weights are computed so as to minimize the estimation variance. The gain in variance of the MVFT with respect to the SFT increases with the number of available signals, as shown by Monte Carlo simulations. Third, when the signal under study contains harmonic oscillations, it makes sense to take into account the underlying structure. The HFT is based on this observation. It estimates the instantaneous fundamental frequency and extract the harmonic components with a weighting procedure similar to the one used in the MVFT.

A Non-Causal Method for Enhancing Frequency Trajectories

When a signal is processed offline, which is not uncommon, all its samples are available at once. In such cases, there is no need to enforce causality in methods for frequency estimation. The NFE takes advantage of this kind of situations in order to achieve better estimation performance by using future information in addition to past information. More specifically, this non-causal iterative procedure refines an initial frequency trajectory by minimizing a cost function derived from the complex discrete oscillator equation. Therefore, once it has converged, the oscillatory behavior of the extracted component is maximized. However, due to its iterative nature, the NFE is much more computationally intensive than the SFT. This may become a problem depending on the considered application.

A Study of Cross-Frequency Coupling Mechanisms in EEG Data

The temporal evolution of phase-amplitude and phase-phase couplings between different frequency bands were analyzed in EEG data recorded during a visual evoked potential experiment with illusory contour stimuli. This study had two primary outcomes. First, the usefulness of the SFT for measuring such cross-frequency interactions was confirmed. Indeed, the adaptive frequency tracking scheme detected more clearly the differences in terms of both phase-amplitude and phase-phase couplings between stimuli with and without illusory contours than a traditional filter-bank approach. This was achieved at the cost of a slight delay. Furthermore, the SFT also highlighted frequency differences in certain bands. Second, the differences in coupling strength followed a complex pattern which depended on the considered frequency bands.

Organization Measures to Predict the Outcome of AF Ablation

The HFT was used to define two organization measures based on adaptive frequency tracking from ECG data. The first one quantifies the power repartition in the frequency domain of AF signals, whereas the second is based on the phase difference between harmonic components. These two measures were developed to assess the AF dynamics and help predicting the outcome of stepwise catheter ablation of long-standing persistent AF. In particular, they indicated that patients who cardioverted during ablation of the left atrium had significantly more organized AF activity.

9.2 Perspectives

Several properties of the adaptive frequency tracking schemes developed in the present thesis are attractive for a wide range of applications. However, they also suffer from some limitations that can become problematic in certain situations. Thus, further developments should focus on these issues. Besides, the outcomes of the EEG and ECG studies presented in this dissertation raise several questions which call for additional investigations. Therefore, a selection of future research directions which would build upon this thesis is given here.

Estimation Delay Compensation

The SFT and its extensions introduce a delay in the estimation of the instantaneous frequency. While limited or even negligible in most cases, this delay can become an issue depending on the dynamics of the signal under study. In particular, after abrupt frequency changes, the algorithms presented above need some time to adapt. Likewise, in high-noise environments, selecting the update parameters so as to achieve sufficient estimation accuracy may result in slow convergence. Therefore, further developments should focus on compensating for the estimation delay.

Detection of Arising and Vanishing Components

The assumption that a frequency component is present for the entire duration of a signal is inaccurate in some cases. Indeed, the instantaneous frequency should be estimated only when an oscillation is active. Furthermore, the corresponding arising and vanishing instants could provide pertinent information about the underlying process. Detecting these instants would be particularly valuable for the extension of the SFT to multiple frequencies, the MFT, for which the number of component present in a signal has to be known *a priori*.

Extensions for Non-causal Frequency Estimation

The NFE can enhance the frequency trajectory of a single oscillation. However, similarly to the SFT, it would be beneficial to extend this iterative method to specific cases, such as the presence of more than one component. While extending the NFE so as to estimate the instantaneous frequency of an oscillation present in several signals is rather straightforward owing to the weighting procedure used in the MVFT, the development of other extensions could lead to difficulties. Furthermore, to remain attractive in practice, these extensions should not increase the computational load beyond reasonable limits.

Investigation of Cross-frequency Couplings in Pathological Cases

The analysis of cross-frequency couplings between neuronal oscillations in different frequency bands during a visual evoked potential experiment showed interesting results. Illusory contour perception elicited a complex pattern of increases and decreases of coupling strength. However, only healthy subjects were included in this investigation. Therefore, future studies should also include patients suffering from different conditions, such as schizophrenia. Contrasting the patterns of couplings obtained with healthy subjects and patients could lead to a better understanding of the pathological mechanisms. It might even help to indicate potential therapeutic solutions.

Classifier for Predicting the Outcome of AF Ablation

Two organization measures derived from the HFT showed significant differences between patients with different AF ablation outcomes. The main limitation of this study was the low number of patients and their poor repartition among the three groups. When enough patients are available, the goal is to develop a classifier to predict reliably AF ablation outcome. Indeed, knowing *a priori* the result of the procedure could help to improve the quality of life of the patients and decrease the management costs of this arrhythmia. Furthermore, the HFT can be used to define additional organization measures. Collectively, these novel measures could also help to determine the outcome of other treatments for AF.

Notation and Definitions



To avoid any confusion or ambiguity, the notation used throughout this dissertation as well as a few essential definitions are provided in this appendix.

A.1 Notation

In order to distinguish continuous- and discrete-time data the following notation is adopted. The temporal variable is indicated with round parentheses (i.e. $x(\cdot)$) for continuous-time signals and with square brackets (i.e. $x[\cdot]$) for discrete-time signals. Moreover, in most cases, the temporal variable is t and n for continuous- and discrete-time data respectively. When working with complex numbers, the imaginary unit is denoted by j . The real and imaginary parts of a complex number are indicated with $\text{Re}\{\cdot\}$ and $\text{Im}\{\cdot\}$. In addition, complex conjugation is denoted with an upper bar.

A.2 Definitions

The definitions of the Fourier transform (FT), the discrete-time Fourier transform (DTFT) and the discrete Fourier transform (DFT) used in this work are given here. The FT, denoted as $\mathcal{F}\{\cdot\}$, and its inverse are expressed as follows,

$$X(f) = \mathcal{F}\{x(t)\} = \int_{-\infty}^{\infty} x(t)e^{-j2\pi ft} dt, \quad (\text{A.1a})$$

$$x(t) = \mathcal{F}^{-1}\{X(f)\} = \int_{-\infty}^{\infty} X(f)e^{j2\pi ft} df. \quad (\text{A.1b})$$

The DTFT is denoted as $\text{DTFT}\{\cdot\}$. It is defined, together with its inverse, as

$$X(e^{j\omega}) = \text{DTFT}\{x[n]\} = \sum_{n=-\infty}^{\infty} x[n]e^{-j\omega n}, \quad (\text{A.2a})$$

$$x[n] = \text{DTFT}^{-1}\{X(e^{j\omega})\} = \frac{1}{2\pi} \int_{-\pi}^{\pi} X(e^{j\omega})e^{j\omega n} d\omega. \quad (\text{A.2b})$$

Finally, the DFT, denoted as $\text{DFT}\{\cdot\}$, and its inverse are given by

$$X[k] = \text{DFT}\{x[n]\} = \sum_{n=0}^{N-1} x[n] e^{-j\frac{2\pi}{N}kn}, \quad k = 0, 1, \dots, N-1, \quad (\text{A.3a})$$

$$x[n] = \text{DFT}^{-1}\{X[k]\} = \frac{1}{N} \sum_{k=0}^{N-1} X[k] e^{j\frac{2\pi}{N}kn}, \quad n = 0, 1, \dots, N-1, \quad (\text{A.3b})$$

where N is the length of $x[n]$ and $X[k]$. At last, two specific functions are introduced: the sign function, defined as

$$\text{sgn}(x) = \begin{cases} -1 & \text{for } x < 0, \\ 0 & \text{for } x = 0, \\ 1 & \text{for } x > 0, \end{cases} \quad (\text{A.4})$$

and the Kronecker delta [8] which is given by

$$\delta[n] = \begin{cases} 1 & \text{for } n = 0, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.5})$$

Performance Analysis of the Single Frequency Tracker

B

This appendix presents the performance analysis of the single frequency tracker (SFT) [64] introduced in Chapter 3. The results given in Section 3.3 are repeated here in more detail, and all calculations are provided. The four different input signals considered for this analysis are described in Section B.1. Then, in Section B.2, the cost function of the adaptive scheme is investigated in detail for these signals. Finally, its frequency estimation performance is quantified in terms of bias and variance in Sections B.3 and B.4, and theoretical values are compared to the outcomes of numerical simulations.

B.1 Input Signals

The four input signals considered for the analysis of the performance of the SFT can be expressed in the form of a generic cisoid embedded in zero-mean noise,

$$y[n] = A_0 e^{j\omega_0 n} + v[n], \quad (\text{B.1})$$

where A_0 is the complex amplitude, ω_0 is the frequency, and $v[n]$ is an additive zero-mean noise process. It is also assumed that the noise is wide-sense stationary [8]. Its other statistical properties, such as the autocorrelation $R_{vv}[k] = E\{v[n]\bar{v}[n-k]\}$, are left undefined to remain as general as possible. They will be specified in the description of each input signal.

First, a cisoid embedded in zero-mean complex white noise is considered since the SFT is based on this specific input. However, the band-pass filter (3.2) is not applied in order to analyze the behavior of the adaptive mechanism solely. The second type is a sinusoid with additive zero-mean real white noise, as it is much more common to work with real data in practice. However, the frequency is estimated from the analytic representation of the input signal, which is computed using the discrete Hilbert transform (DHT). For the same reasons as for the first signal, no filtering is used. The third and fourth input signals are the same as the two first ones, except that the band-pass filter (3.2) is applied before performing the estimation. The transfer function of the filter is repeated here in a slightly modified form:

$$G(e^{j\omega}) = \frac{1 - \beta}{1 - \beta e^{j\omega_c} e^{-j\omega}}, \quad (\text{B.2})$$

where ω_c is the central frequency and β ($0 \ll \beta < 1$) controls the bandwidth, as in the description of the SFT. The four input signals are summarized as follows:

Noisy cisoid A cisoid embedded in additive zero-mean complex white noise.

Noisy sinusoid The analytic representation of a sinusoid embedded in additive zero-mean real white noise.

Filtered noisy cisoid A cisoid embedded in additive zero-mean complex white noise filtered with the band-pass filter (B.2).

Filtered noisy sinusoid The analytic representation of a sinusoid embedded in additive zero-mean real white noise filtered with the band-pass filter (B.2).

By means of these four different input signals, the effects of real data and band-pass filtering can be investigated separately. Also, it should be mentioned that the first two signals correspond to a simplified SFT without band-pass filtering. The following sections present these signals in detail alongside the autocorrelation of their respective noise process.

B.1.1 Noisy Cisoid

The cisoid embedded in additive zero-mean complex white noise is defined as

$$y_1[n] = A_0 e^{j\omega_0 n} + v_1[n], \quad (\text{B.3})$$

where A_0 and ω_0 are respectively the complex amplitude and frequency, and $v_1[n]$ is the complex noise process. As the noise is white, its autocorrelation is

$$R_{v_1 v_1}[k] = \sigma^2 \delta[k], \quad (\text{B.4})$$

with σ^2 the variance and $\delta[k]$ the Kronecker delta (A.5). In addition, the real and imaginary parts of the noise are assumed to be independent and to have equal variances (i.e. $\text{Var}\{\text{Re}\{v_1[n]\}\} = \text{Var}\{\text{Im}\{v_1[n]\}\} = \sigma^2/2$).

B.1.2 Noisy Sinusoid

The sinusoid with additive zero-mean real white noise is given by

$$x_2[n] = B_0 \sin(\omega_0 n + \phi_0) + w_2[n], \quad (\text{B.5})$$

where B_0 , ω_0 and ϕ_0 are respectively the amplitude, frequency and phase, and $w_2[n]$ is the real noise process with autocorrelation $R_{w_2 w_2}[k] = \sigma^2 \delta[k]$. Using the fact that the analytic representation of a sinusoid is a cisoid, the input signal is

$$\begin{aligned} y_2[n] &= x_2[n] + jh[n] * x_2[n] \\ &= B_0 e^{j\phi_0} e^{j\omega_0 n} + w_2[n] + jw_{2,h}[n] \\ &= A_0 e^{j\omega_0 n} + v_2[n], \end{aligned} \quad (\text{B.6})$$

where $h[n]$ is the Hilbert filter (2.12), $A_0 = B_0 e^{j\phi_0}$ collects the amplitude and phase, $w_{2,h}[n] = h[n] * w_2[n]$, and $v_2[n] = w_2[n] + jw_{2,h}[n]$. In the previous expressions, the notation $*$ denotes the convolution operation. The autocorrelation of the analytic representation of the noise is obtained as follows:

$$\begin{aligned} R_{v_2 v_2}[k] &= \text{E} \left\{ (w_2[n] + jw_{2,h}[n])(w_2[n-k] - jw_{2,h}[n-k]) \right\} \\ &= R_{w_2 w_2}[k] - jR_{w_2 w_{2,h}}[k] + jR_{w_{2,h} w_2}[k] + R_{w_{2,h} w_{2,h}}[k]. \end{aligned} \quad (\text{B.7})$$

Since $w_2[n]$ is wide-sense stationary and $h[n]$ is a linear shift-invariant filter, $R_{w_2,h w_2,h}[k]$, $R_{w_2 w_2,h}[k]$ and $R_{w_2,h w_2}[k]$ can be computed via the power spectral densities [3, 200]:

$$S_{w_2,h w_2,h}(e^{j\omega}) = |H(e^{j\omega})|^2 \cdot S_{w_2 w_2}(e^{j\omega}) = \sigma^2 |H(e^{j\omega})|^2 \xrightarrow{\text{DTFT}^{-1}} R_{w_2,h w_2,h}[k] = \sigma^2 \delta[k], \quad (\text{B.8a})$$

$$S_{w_2 w_2,h}(e^{j\omega}) = \bar{H}(e^{j\omega}) \cdot S_{w_2 w_2}(e^{j\omega}) = \sigma^2 \bar{H}(e^{j\omega}) \xrightarrow{\text{DTFT}^{-1}} R_{w_2 w_2,h}[k] = -\sigma^2 h[k], \quad (\text{B.8b})$$

$$S_{w_2,h w_2}(e^{j\omega}) = H(e^{j\omega}) \cdot S_{w_2 w_2}(e^{j\omega}) = \sigma^2 H(e^{j\omega}) \xrightarrow{\text{DTFT}^{-1}} R_{w_2,h w_2}[k] = \sigma^2 h[k], \quad (\text{B.8c})$$

where $H(e^{j\omega})$ is the discrete-time Fourier transform (DTFT) of the Hilbert filter (2.14). Taken together, these results yield the autocorrelation of the analytic representation of the input noise:

$$R_{v_2 v_2}[k] = 2\sigma^2(\delta[k] + jh[k]). \quad (\text{B.9})$$

Therefore, although $w_2[n]$ is white, its analytic representation is not.

B.1.3 Filtered Noisy Cisoid

Similarly to the first input signal, the noisy cisoid before the filtering operation is defined as

$$x_3[n] = A_0 e^{j\omega_0 n} + w_3[n], \quad (\text{B.10})$$

where A_0 and ω_0 are respectively the complex amplitude and frequency, and $w_3[n]$ is a zero-mean complex white noise with autocorrelation $R_{w_3 w_3}[k] = \sigma^2 \delta[k]$. Similar to the noisy cisoid case, the variances of the real and imaginary parts of the noise are assumed to be independent and identical. The filtered input signal is obtained as follows:

$$\begin{aligned} y_3[n] &= g[n] * (A_0 e^{j\omega_0 n} + w_3[n]) \\ &= \frac{1 - \beta}{1 - \beta e^{-j(\omega_0 - \omega_c)}} A_0 e^{j\omega_0 n} + g[n] * w_3[n] \\ &= \frac{1 - \beta}{1 - \beta e^{-j(\omega_0 - \omega_c)}} A_0 e^{j\omega_0 n} + v_3[n], \end{aligned} \quad (\text{B.11})$$

where $g[n]$ is the impulse response of (B.2) and $v_3[n] = g[n] * w_3[n]$. So, depending on the central frequency, the filter changes the amplitude as well as the phase offset of the cisoid. However, when ω_c coincides with ω_0 , the cisoid is not altered. On the other hand, the filtered noise is no more white. Its autocorrelation is computed through the power spectral density by transforming $R_{w_3 w_3}[k]$ to the frequency domain:

$$S_{v_3 v_3}(e^{j\omega}) = |G(e^{j\omega})|^2 \cdot S_{w_3 w_3}(e^{j\omega}) = \frac{(1 - \beta)^2}{1 - 2\beta \cos(\omega - \omega_c) + \beta^2} \cdot \sigma^2. \quad (\text{B.12})$$

Reverting back to the time domain yields the autocorrelation of the filtered noise,

$$R_{v_3 v_3}[k] = \sigma^2 \frac{1 - \beta}{1 + \beta} \beta^{|k|} e^{j\omega_c k}. \quad (\text{B.13})$$

B.1.4 Filtered Noisy Sinusoid

Before computing its analytic representation and applying the band-pass filter, the fourth input signal is

$$x_4[n] = B_0 \sin(\omega_0 n + \phi_0) + w_4[n], \quad (\text{B.14})$$

where B_0 , ω_0 and ϕ_0 are respectively the amplitude, frequency and phase, and $w_4[n]$ is a zero-mean real white noise with autocorrelation $R_{w_4 w_4}[k] = \sigma^2 \delta[k]$. The filtered analytic representation of the previous expression is computed as

$$\begin{aligned} y_4[n] &= g[n] * (x_4[n] + jh[n] * x_4[n]) \\ &= g[n] * (B_0 e^{j\phi_0} e^{j\omega_0 n} + w_4[n] + jw_{4,h}[n]) \\ &= g[n] * A_0 e^{j\omega_0 n} + g[n] * (w_4[n] + jw_{4,h}[n]) \\ &= \frac{1 - \beta}{1 - \beta e^{-j(\omega_0 - \omega_c)}} A_0 e^{j\omega_0 n} + v_4[n], \end{aligned} \quad (\text{B.15})$$

where $g[n]$ is the impulse response of (B.2), $h[n]$ is the Hilbert filter (2.12), $A_0 = B_0 e^{j\phi_0}$ collects the amplitude and phase, $w_{4,h}[n] = h[n] * w_4[n]$, and $v_4[n] = g[n] * (w_4[n] + jw_{4,h}[n])$. Using the result for the autocorrelation of $v_2[n]$ (B.9), the power spectral density of the noise $v_4[n]$ is given by the following expression:

$$\begin{aligned} S_{v_4 v_4}(e^{j\omega}) &= |G(e^{j\omega})|^2 \cdot (S_{w_4 w_4}(e^{j\omega}) - jS_{w_4 w_{4,h}}(e^{j\omega}) + jS_{w_{4,h} w_4}(e^{j\omega}) + S_{w_{4,h} w_{4,h}}(e^{j\omega})) \\ &= |G(e^{j\omega})|^2 \cdot \sigma^2 (2 + jH(e^{j\omega}) - j\bar{H}(e^{j\omega})). \end{aligned} \quad (\text{B.16})$$

Transforming back to the time domain, the autocorrelation of $v_4[n]$ is

$$\begin{aligned} R_{v_4 v_4}[k] &= \left(\frac{1 - \beta}{1 + \beta} \beta^{|k|} e^{j\omega_c k} \right) * (2\sigma^2 (\delta[k] + jh[k])) \\ &= 2\sigma^2 \frac{1 - \beta}{1 + \beta} ((\delta[k] + jh[k]) * \beta^{|k|} e^{j\omega_c k}) = 2\sigma^2 \frac{1 - \beta}{1 + \beta} a[k], \end{aligned} \quad (\text{B.17})$$

with

$$\begin{aligned} a[k] &= (\delta[k] + jh[k]) * \beta^{|k|} e^{j\omega_c k} \\ &= \sum_{l=-\infty}^{+\infty} (\delta[l] + jh[l]) \cdot \beta^{|k-l|} e^{j\omega_c (k-l)} \\ &= \beta^{|k|} e^{j\omega_c k} + j \sum_{l=1}^{+\infty} h[l] \beta^{k-l} e^{j\omega_c (k-l)} - j \sum_{l=1}^{+\infty} h[l] \beta^{k+l} e^{j\omega_c (k+l)} \\ &= \beta^{|k|} e^{j\omega_c k} + \frac{2j}{\pi} e^{j\omega_c k} \sum_{l=0}^{+\infty} \frac{\beta^{k-2l-1}}{2l+1} e^{-j\omega_c (2l+1)} - \frac{2j}{\pi} e^{j\omega_c k} \sum_{l=0}^{+\infty} \frac{\beta^{k+2l+1}}{2l+1} e^{j\omega_c (2l+1)}. \end{aligned} \quad (\text{B.18})$$

Using the Taylor series of $\ln((1+r)/(1-r))$ around $r=0$ [201],

$$\ln\left(\frac{1+r}{1-r}\right) = 2 \sum_{l=0}^{+\infty} \frac{r^{2l+1}}{2l+1}, \quad |r| < 1, \quad (\text{B.19})$$

$a[k]$ can be rewritten as follows:

$$a[k] = \beta^{|k|} e^{j\omega_c k} - \frac{j}{\pi} e^{j\omega_c k} \left(\beta^k \ln\left(\frac{1 + \beta e^{j\omega_c}}{1 - \beta e^{j\omega_c}}\right) - \beta^{-k} \ln\left(\frac{1 + \beta e^{-j\omega_c}}{1 - \beta e^{-j\omega_c}}\right) \right) + \frac{2j}{\pi} e^{j\omega_c k} b[k], \quad (\text{B.20})$$

where $b[k]$ is defined as

$$b[k] = \begin{cases} 0 & \text{for } k = -1, 0, 1, \\ \beta^k \sum_{l=0}^{\lfloor k/2 \rfloor - 1} \frac{(\beta^{-1} e^{-j\omega_c})^{2l+1}}{2l+1} - \beta^{-k} \sum_{l=0}^{\lfloor k/2 \rfloor - 1} \frac{(\beta e^{-j\omega_c})^{2l+1}}{2l+1} & \text{for } k > 1, \\ \bar{b}[-k] & \text{for } k < -1. \end{cases} \quad (\text{B.21})$$

In the previous expression, the notation $\lfloor \cdot \rfloor$ denotes the integer part (or floor function). Therefore, using (B.17), (B.20) and (B.21), the noise autocorrelation $R_{v_4 v_4}[k]$ is now given by

$$R_{v_4 v_4}[k] = 2\sigma^2 \frac{1-\beta}{1+\beta} e^{j\omega_c k} \left(\beta^{|k|} - \frac{j}{\pi} \left(\beta^k \ln \left(\frac{1+\beta e^{j\omega_c}}{1-\beta e^{j\omega_c}} \right) - \beta^{-k} \ln \left(\frac{1+\beta e^{-j\omega_c}}{1-\beta e^{-j\omega_c}} \right) \right) + \frac{2j}{\pi} b[k] \right). \quad (\text{B.22})$$

B.2 Cost Function Analysis

The cost function (3.6) of the SFT, which is based on the complex discrete oscillator equation, is defined as

$$J(\omega) = \text{E} \left\{ |y[n] - e^{j\omega} y[n-1]|^2 \right\}. \quad (\text{B.23})$$

In order to analyze this cost function, the first step is to derive its expression for the generic noisy cisoid $y[n]$ (B.1). Thus, inserting (B.1) into (B.23) yields

$$\begin{aligned} J(\omega) &= \text{E} \left\{ |A_0 e^{j\omega_0 n} + v[n] - A_0 e^{j(\omega_0(n-1)+\omega)} - e^{j\omega} v[n-1]|^2 \right\} \\ &= 2|A_0|^2 (1 - \cos(\omega - \omega_0)) + 2(\text{Re}\{R_{vv}[0] - \text{Re}\{e^{-j\omega} R_{vv}[1]\}\}). \end{aligned} \quad (\text{B.24})$$

Finally, using trigonometric identities [202], the cost function reduces to

$$J(\omega) = 4|A_0|^2 \sin^2 \left(\frac{\omega - \omega_0}{2} \right) + 2(\text{Re}\{R_{vv}[0] - \text{Re}\{e^{-j\omega} R_{vv}[1]\}\}). \quad (\text{B.25})$$

The first term is minimized for $\omega = \omega_0$. Hence, as long as the second term does not depend on ω , minimizing this cost function leads to an unbiased estimate. The cost function for each considered input signal is easily computed by replacing $R_{vv}[k]$ with the corresponding noise autocorrelation.

B.2.1 Noisy Cisoid

Using the expression for $R_{v_1 v_1}[k]$ (B.4), the cost function (B.25) becomes

$$J(\omega) = 4|A_0|^2 \sin^2 \left(\frac{\omega - \omega_0}{2} \right) + 2\sigma^2. \quad (\text{B.26})$$

Consequently, the frequency estimate obtained with this cost function is unbiased as shown in Figure B.1 for different noise variances. This situation corresponds to a simplified version of the SFT since the filtering operation is missing.

B.2.2 Noisy Sinusoid

Inserting the autocorrelation $R_{v_2 v_2}[k]$ (B.9) into the cost function (B.25) leads to

$$J(\omega) = 4|A_0|^2 \sin^2 \left(\frac{\omega - \omega_0}{2} \right) + 4\sigma^2 \left(1 - \frac{2}{\pi} \sin(\omega) \right). \quad (\text{B.27})$$

In this case, minimizing the cost function results in a biased estimate. More precisely, it is biased towards the frequency $\omega = 0.5\pi$ since the second term is minimized for this frequency. Figure B.2 illustrates the bias for different noise variances. As for the noisy cisoid, the band-pass filter (B.2) is missing. Ultimately, the bias is caused by the DHT which introduces long-term correlations in the analytic representation of the noise.

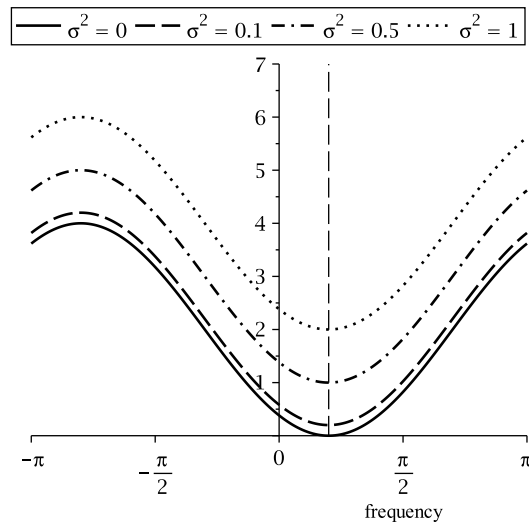


Figure B.1: Oscillator-based cost function for a noisy cisoid. The cost function is plotted for four different noise variances. The amplitude and frequency are set to $A_0 = 1$ and $\omega_0 = 0.2\pi$. The vertical dashed line denotes the true frequency of the cisoid.

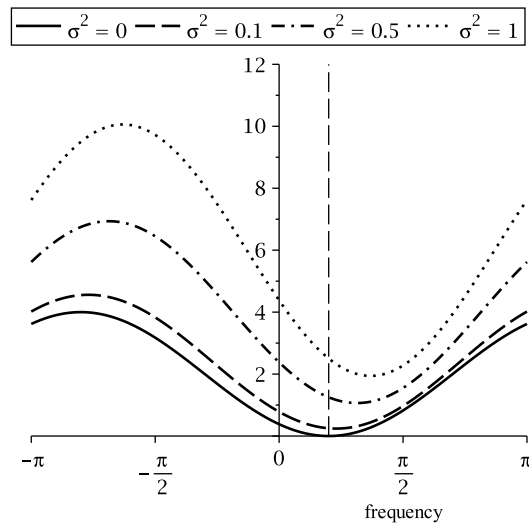


Figure B.2: Oscillator-based cost function for the analytic representation of a noisy sinusoid. The cost function is plotted for four different noise variances. The amplitude and frequency are set to $A_0 = 1$ and $\omega_0 = 0.2\pi$. The vertical dashed line denotes the true frequency of the sinusoid.

B.2.3 Filtered Noisy Cisoid

Replacing $R_{vv}[k]$ in (B.25) by $R_{v_3v_3}[k]$ (B.13) yields the following cost function:

$$J(\omega) = \frac{(1-\beta)^2}{1-2\beta\cos(\omega_0-\omega_c)+\beta^2} \cdot 4|A_0|^2 \sin^2\left(\frac{\omega-\omega_0}{2}\right) + 2\sigma^2 \frac{1-\beta}{1+\beta} (1-\beta\cos(\omega-\omega_c)). \quad (\text{B.28})$$

Therefore, the tracking algorithm is biased towards the central frequency ω_c of the band-pass filter for this input signal. However, when the central frequency matches the frequency of the cisoid ($\omega_c = \omega_0$), the cost function simplifies into

$$J(\omega) = 4|A_0|^2 \sin^2\left(\frac{\omega-\omega_0}{2}\right) + 2\sigma^2 \frac{1-\beta}{1+\beta} (1-\beta\cos(\omega-\omega_0)), \quad (\text{B.29})$$

and the frequency estimate is unbiased. The cost function for this type of input signal is plotted in Figure B.3 for different noise variances and central frequencies. This case corresponds to the true SFT compared to the two previous signals since the input signal is band-pass filtered. The important result is that, when the true frequency and the central frequency are in close agreement, the bias is very small. It even vanishes when they coincide. The bandwidth of the filter (B.2) also plays a significant role. Indeed, narrow bandwidths tend to increase the bias. However, as long as the frequency estimate provided by the SFT is close to the true frequency and the bandwidth is not exaggeratedly narrow, the bias should remain negligible.

B.2.4 Filtered Noisy Sinusoid

Using $R_{v_4v_4}[0]$ and $R_{v_4v_4}[1]$ (B.22), the cost function (B.25) for a filtered noisy sinusoid becomes

$$\begin{aligned} J(\omega) &= \frac{(1-\beta)^2}{1-2\beta\cos(\omega_0-\omega_c)+\beta^2} \cdot 4|A_0|^2 \sin^2\left(\frac{\omega-\omega_0}{2}\right) \\ &\quad + 4\sigma^2 \frac{1-\beta}{1+\beta} \cdot (a[0] - \text{Re}\{e^{-j\omega} a[1]\}) \\ &= \frac{(1-\beta)^2}{1-2\beta\cos(\omega_0-\omega_c)+\beta^2} \cdot 4|A_0|^2 \sin^2\left(\frac{\omega-\omega_0}{2}\right) \\ &\quad + 4\sigma^2 \frac{1-\beta}{1+\beta} \cdot \left(1 + \frac{j}{\pi} \ln\left(\frac{1-\beta^2-2j\beta\sin(\omega_c)}{1-\beta^2+2j\beta\sin(\omega_c)}\right)\right. \\ &\quad \left. - \text{Re}\left\{\beta e^{-j(\omega-\omega_c)} + \frac{j e^{-j(\omega-\omega_c)}}{\pi\beta} \left(\ln\left(\frac{1+\beta e^{-j\omega_c}}{1-\beta e^{-j\omega_c}}\right) - \beta^2 \ln\left(\frac{1+\beta e^{j\omega_c}}{1-\beta e^{j\omega_c}}\right)\right)\right\}\right) \quad (\text{B.30}) \\ &= \frac{(1-\beta)^2}{1-2\beta\cos(\omega_0-\omega_c)+\beta^2} \cdot 4|A_0|^2 \sin^2\left(\frac{\omega-\omega_0}{2}\right) \\ &\quad + 4\sigma^2 \frac{1-\beta}{1+\beta} \cdot \left(1 - \beta\cos(\omega-\omega_c) + \frac{j}{\pi} \ln\left(\frac{1-\beta^2-2j\beta\sin(\omega_c)}{1-\beta^2+2j\beta\sin(\omega_c)}\right)\right. \\ &\quad \left. + \text{Im}\left\{\frac{e^{-j(\omega-\omega_c)}}{\pi\beta} \left(\ln\left(\frac{1+\beta e^{-j\omega_c}}{1-\beta e^{-j\omega_c}}\right) - \beta^2 \ln\left(\frac{1+\beta e^{j\omega_c}}{1-\beta e^{j\omega_c}}\right)\right)\right\}\right). \end{aligned}$$

Minimizing this expression yields a biased estimate of the frequency. Two different factors are responsible for the bias: the DHT and the band-pass filter (B.2). Nevertheless, when the central frequency of the filter is aligned with the true frequency of the cisoid ($\omega_c = \omega_0$), the bias due to

the filtering operation disappears and the cost function reduces to

$$\begin{aligned}
J(\omega) = & 4|A_0|^2 \sin^2\left(\frac{\omega - \omega_0}{2}\right) \\
& + 4\sigma^2 \frac{1 - \beta}{1 + \beta} \cdot \left(1 - \beta \cos(\omega - \omega_0) + \frac{j}{\pi} \ln\left(\frac{1 - \beta^2 - 2j\beta \sin(\omega_0)}{1 - \beta^2 + 2j\beta \sin(\omega_0)}\right)\right) \\
& + \text{Im} \left\{ \frac{e^{-j(\omega - \omega_0)}}{\pi\beta} \left(\ln\left(\frac{1 + \beta e^{-j\omega_0}}{1 - \beta e^{-j\omega_0}}\right) - \beta^2 \ln\left(\frac{1 + \beta e^{j\omega_0}}{1 - \beta e^{j\omega_0}}\right) \right) \right\}.
\end{aligned} \tag{B.31}$$

Figure B.4 shows this cost function for different combinations of noise variances and central frequencies. The important point is that, for a sufficiently narrow bandwidth and when the central frequency of the filter coincides with the frequency of the cisoid, the bias due to the DHT becomes negligible.

B.3 Bias Analysis

Minimizing the cost function (B.25) can result in a biased frequency estimate under certain circumstances, as discussed in the previous section. The dependence of the bias on the various parameters of the input signal and of the SFT remains unclear. Therefore, it is now analyzed for the four considered input signals. In fact, for the sake of clarity, an approximation of the bias is analyzed. Indeed, the SFT is designed to estimate the autocorrelation at lag one of the signal $y[n]$, $R_{yy}[1]$, with an exponentially weighted average (3.9). The instantaneous frequency estimate at each time is then obtained by taking the argument of this average: $\omega[n + 1] = \arg\{Q[n]\}$. Consequently, an approximation of the expected value of the frequency estimate is given by

$$\tilde{\omega} = \arg\{E\{Q[n]\}\}, \tag{B.32}$$

and the approximated bias is defined as follows:

$$\text{Bias}\{\tilde{\omega}\} = \tilde{\omega} - \omega_0 = \arg\{E\{Q[n]\}\} - \omega_0, \tag{B.33}$$

where ω_0 is the true frequency.

In order to evaluate this approximation, the first step is to compute the expected value of $Q[n]$, $\tilde{Q} = E\{Q[n]\}$. Taking the expectation of the recursion (3.9) yields the following expression,

$$\tilde{Q} = E\{Q[n]\} = \delta E\{Q[n - 1]\} + (1 - \delta) E\{y[n]\bar{y}[n - 1]\} = \delta\tilde{Q} + (1 - \delta)R_{yy}[1], \tag{B.34}$$

which in turn leads to $\tilde{Q} = R_{yy}[1]$. Therefore, when the input signal is the generic cisoid embedded in additive zero-mean noise (B.1), $y[n] = A_0 e^{j\omega_0 n} + v[n]$, the expected value of $Q[n]$ is given by

$$\tilde{Q} = |A_0|^2 e^{j\omega_0} + R_{vv}[1]. \tag{B.35}$$

It follows that the approximation of the expected value of the frequency estimate for the generic noisy cisoid is

$$\tilde{\omega} = \arg\{\tilde{Q}\} = \arg\{|A_0|^2 e^{j\omega_0} + R_{vv}[1]\}, \tag{B.36}$$

and the approximated bias is

$$\text{Bias}\{\tilde{\omega}\} = \arg\{|A_0|^2 e^{j\omega_0} + R_{vv}[1]\} - \omega_0. \tag{B.37}$$

Thus, as long as $R_{vv}[1]$ is zero or its argument equals ω_0 , the SFT provides an unbiased estimate of the frequency.

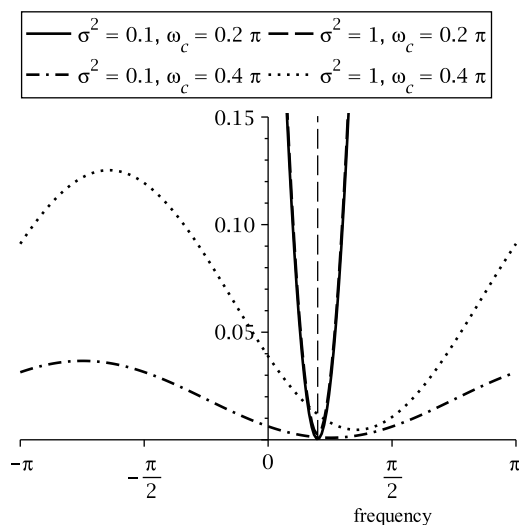


Figure B.3: Oscillator-based cost function for a filtered noisy cisoid. The cost function is plotted for four combinations of noise variances and central frequencies. The amplitude and frequency are set to $A_0 = 1$ and $\omega_0 = 0.2\pi$, and the bandwidth parameter of the filter (B.2) is set to $\beta = 0.95$. The vertical dashed line denotes the true frequency of the cisoid.

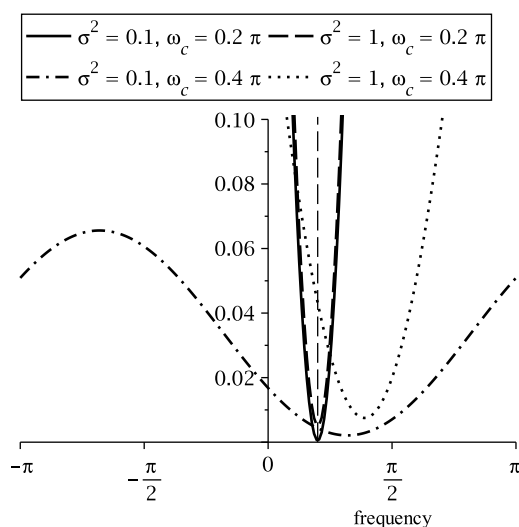


Figure B.4: Oscillator-based cost function for the filtered analytic representation of a noisy sinusoid. The cost function is plotted for four combinations of noise variances and central frequencies. The amplitude and frequency are set to $A_0 = 1$ and $\omega_0 = 0.2\pi$, and the bandwidth parameter of the filter (B.2) is set to $\beta = 0.95$. The vertical dashed line denotes the true frequency of the sinusoid.

In addition to theoretical analyses, the bias was also investigated with Monte Carlo simulations for the four input signals described previously: noisy cisoid, noisy sinusoid, filtered noisy cisoid and filtered noisy sinusoid. For the first two cases, a simplified SFT without the adaptive filtering operation was applied, while for the last two cases, the complete algorithm was used. The input signals were cisoids with uniformly distributed random phase embedded in complex circular Gaussian white noise for the noisy cisoid and filtered noisy cisoid cases. Whereas, for the two sinusoids cases, the input signals were sinusoids with uniformly distributed random phase with additive real Gaussian white noise. The bias was computed over the last 1000 samples of 2000 iterations and averaged over 10,000 runs. This simulated bias was then compared to the theoretical one. It should be noted that, when not specified differently, the following parameter values were selected: the amplitude and frequency of the cisoids and sinusoids were set to $A_0 = 1$ and $\omega_0 = 0.4\pi$, the SNR was set to 10 dB, the forgetting factor of the SFT was set to $\delta = 0.95$, the bandwidth parameter of the adaptive band-pass filter (B.2) was set to $\beta = 0.95$ and its central frequency was given by the current frequency estimate of the SFT (3.10). The parameters of the adaptive band-pass filter are only relevant for the last two cases where the complete SFT was applied.

B.3.1 Noisy Cisoid

Similar to the cost function analysis, $R_{v_1 v_1}[1]$ (B.4) is plugged into the general expression for the bias (B.37). This gives an unbiased estimate of the instantaneous frequency:

$$\text{Bias}\{\tilde{\omega}\} = \arg\left\{|A_0|^2 e^{j\omega_0}\right\} - \omega_0 = 0. \quad (\text{B.38})$$

This result was confirmed with Monte Carlo simulations for SNR values taken from 0 to 20 dB in 1 dB steps. Figure B.5 shows that the simplified SFT is indeed unbiased for a cisoid embedded in white noise.

B.3.2 Noisy Sinusoid

Inserting $R_{v_2 v_2}[1] = j4\sigma^2/\pi$ into (B.37) yields the following expression for the bias.

$$\text{Bias}\{\tilde{\omega}\} = \arg\left\{|A_0|^2 e^{j\omega_0} + j\frac{4\sigma^2}{\pi}\right\} - \omega_0. \quad (\text{B.39})$$

The bias estimate obtained with the Monte Carlo procedure matched closely the theoretical value as shown in Figure B.6 for SNR values between 0 and 20 dB (Figure B.6a) as well as for sinusoid frequencies ranging from 0.04π to 0.96π in 0.04π steps (Figure B.6b). As indicated by the cost function analysis, for this input signal, the tracking algorithm is biased towards 0.5π since the analytic representation of a real white noise is no longer white. Besides, the bias absolute value depended on both the noise level and the distance between the sinusoid frequency and 0.5π .

B.3.3 Filtered Noisy Cisoid

In this case, by replacing $R_{vv}[k]$ with (B.13) in (B.37), the bias of the frequency estimate becomes

$$\text{Bias}\{\tilde{\omega}\} = \arg\left\{\frac{(1-\beta)^2}{1-2\beta\cos(\omega_0-\omega_c)+\beta^2} \cdot |A_0|^2 e^{j\omega_0} + \sigma^2 \frac{1-\beta}{1+\beta} \beta e^{j\omega_c}\right\} - \omega_0. \quad (\text{B.40})$$

Based on this expression, the SFT is unbiased when the central frequency of the band-pass filter is equal to the frequency of the cisoid, $\omega_c = \omega_0$. This is illustrated in Figure B.7a for SNR values taken from 0 to 20 dB in 1 dB steps. However, when the filter is not centered on the cisoid

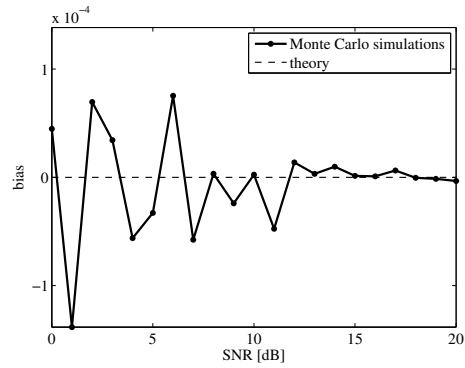


Figure B.5: Bias of the SFT for a noisy cisoid with respect to the SNR ($A_0 = 1$, $\omega_0 = 0.4\pi$, $\delta = 0.95$).

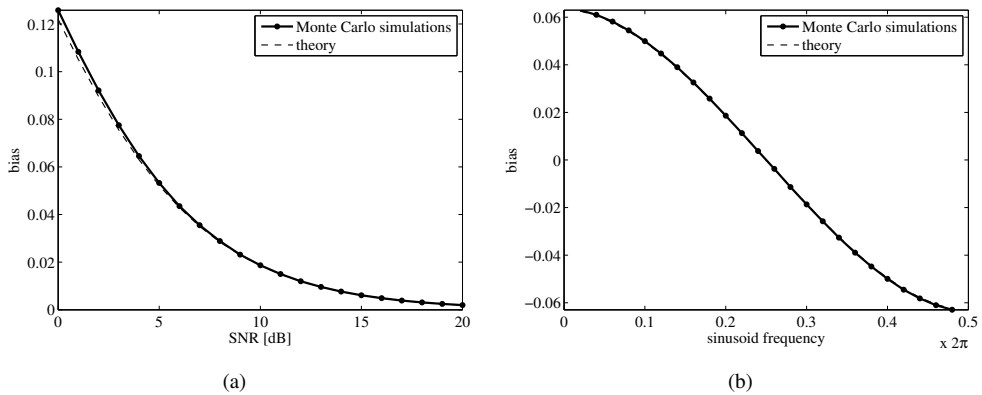


Figure B.6: Bias of the SFT for the analytic representation of a noisy sinusoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\delta = 0.95$) with respect to (a) the SNR, and (b) to the sinusoid frequency ω_0 .

frequency, the tracking algorithm is biased. Monte Carlo simulations were performed in order to evaluate the deviation from the true frequency. For this purpose, fixed central frequencies were taken uniformly between 0.04π and 0.96π in 0.04π steps for the filter (B.2). It is important to clarify that in this particular case the central frequency of the filter was not adapted. The outcomes of the numerical simulations are shown in Figure B.7b alongside theoretical values. From these results, it is clear that the bias is almost negligible when the central frequency is close to the cisoid frequency.

B.3.4 Filtered Noisy Sinusoid

As before, the noise autocorrelation $R_{vv}[k]$ of the generic cisoid is substituted with the one for this input signal (B.22). Thus, the expected value of $Q[n]$ turns into

$$\begin{aligned}\tilde{Q} &= \frac{(1-\beta)^2}{1-2\beta\cos(\omega_0-\omega_c)+\beta^2} \cdot |A_0|^2 e^{j\omega_0} + 2\sigma^2 \frac{1-\beta}{1+\beta} a[1] \\ &= \frac{(1-\beta)^2}{1-2\beta\cos(\omega_0-\omega_c)+\beta^2} \cdot |A_0|^2 e^{j\omega_0} \\ &\quad + 2\sigma^2 \frac{1-\beta}{1+\beta} \left(\beta e^{j\omega_c} + \frac{j e^{j\omega_c}}{\pi\beta} \left(\ln \left(\frac{1+\beta e^{-j\omega_c}}{1-\beta e^{-j\omega_c}} \right) - \beta^2 \ln \left(\frac{1+\beta e^{j\omega_c}}{1-\beta e^{j\omega_c}} \right) \right) \right),\end{aligned}\tag{B.41}$$

and the estimation bias is

$$\text{Bias}\{\tilde{\omega}\} = \arg\{\tilde{Q}\} - \omega_0.\tag{B.42}$$

Like for the previous cases, this theoretical value was compared to the bias computed with Monte Carlo simulations. In particular, the influences of the noise level, the sinusoid frequency ω_0 , the central frequency ω_c and the bandwidth parameter β were analyzed in detail. Figures B.8a and B.8b show the bias of the SFT for SNR values taken from 0 to 20 dB in 1 dB steps and for sinusoid frequencies ω_0 ranging from 0.04π to 0.96π respectively. The effect of the central frequency ω_c on the bias is illustrated in Figure B.8c for fixed values between 0.04π and 0.96π , while the influence of the bandwidth parameter β is shown in Figure B.8d for values taken from 0.5 to 0.95 in 0.05 steps. For all these situations, theoretical and practical bias were in close agreement. Also, since the analytic representation of a white noise is no longer white, the SFT is biased (with the exception of the special case $\omega_0 = \omega_c = 0.5\pi$). Nevertheless, the bias is very limited or even negligible when the band-pass filter is narrow and centered on the sinusoid frequency.

B.4 Variance Analysis

As it was done for the bias in Section B.3, the influence of the SFT parameters on the variance of the instantaneous frequency estimate $\omega[n]$ is analyzed for the four different input signals. However, theoretical values, even approximate ones, are extremely difficult to obtain. In fact, the calculations quickly become intractable due to the argument in the update of the instantaneous frequency estimate (3.10). Nonetheless, it is still possible to derive an expression for the variance of the internal variable $Q[n]$, $\text{Var}\{Q[n]\}$. The computation is based on the update (3.9) which can be rewritten as follows,

$$Q[n] = \delta Q[n-1] + (1-\delta)y[n]\bar{y}[n-1],\tag{B.43}$$

$$Q[n] - \tilde{Q} = \delta(Q[n-1] - \tilde{Q}) + (1-\delta)(y[n]\bar{y}[n-1] - \tilde{Q}),\tag{B.44}$$

$$Q_c[n] = \delta Q_c[n] + (1-\delta)(y[n]\bar{y}[n-1] - R_{yy}[1]),\tag{B.45}$$

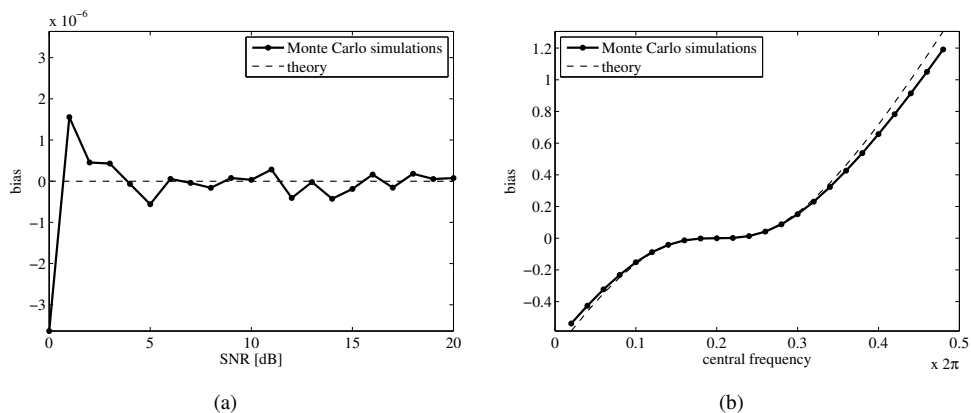


Figure B.7: Bias of the SFT for a filtered noisy cisoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$) with respect (a) to the SNR, and (b) to the central frequency ω_c .

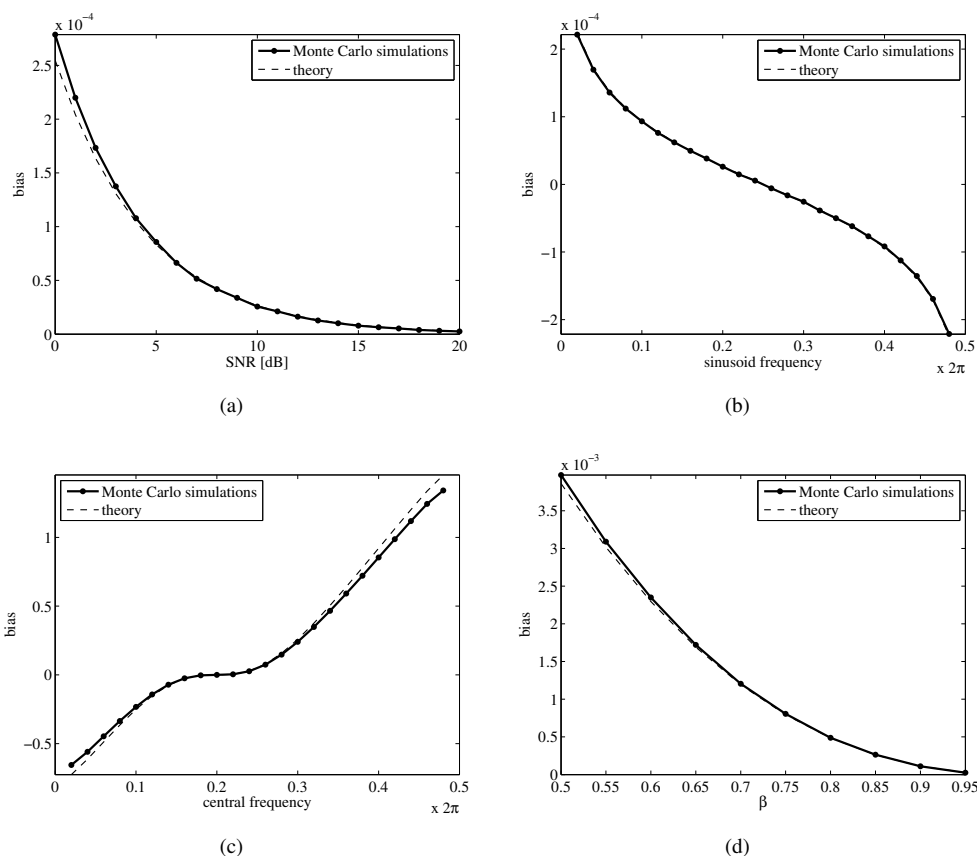


Figure B.8: Bias of the SFT for the filtered analytic representation of a noisy sinusoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$) with respect (a) to the SNR, (b) to the sinusoid frequency ω_0 , (c) to the central frequency ω_c , and (d) to the bandwidth parameter β .

with $Q_c[n] = Q[n] - \bar{Q}$. Taking the squared modulus and then the expectation on both side of the last equation yields

$$\begin{aligned} \mathbb{E} \left\{ |Q_c[n]|^2 \right\} &= \delta^2 \mathbb{E} \left\{ |Q_c[n-1]|^2 \right\} + (1-\delta)^2 \mathbb{E} \left\{ |y[n]\bar{y}[n-1] - R_{yy}[1]|^2 \right\} \\ &\quad + 2\delta(1-\delta) \operatorname{Re} \left\{ \mathbb{E} \left\{ Q_c[n-1] (\bar{y}[n]y[n-1] - \bar{R}_{yy}[1]) \right\} \right\}. \end{aligned} \quad (\text{B.46})$$

Therefore, the variance of $Q[n]$ is given by

$$\begin{aligned} \operatorname{Var}\{Q[n]\} &= R_{Q_c Q_c}[0] = \frac{1-\delta}{1+\delta} \mathbb{E} \left\{ |y[n]\bar{y}[n-1] - R_{yy}[1]|^2 \right\} \\ &\quad + \frac{2\delta}{1+\delta} \operatorname{Re} \left\{ \mathbb{E} \left\{ Q_c[n-1] (\bar{y}[n]y[n-1] - \bar{R}_{yy}[1]) \right\} \right\}. \end{aligned} \quad (\text{B.47})$$

When the input signal is the generic cisoid with additive zero-mean white noise (B.1), $y[n] = A_0 e^{j\omega_0 n} + v[n]$, the two expectations in the right-hand side of the previous equation can be rewritten as

$$\begin{aligned} \mathbb{E} \left\{ |y[n]\bar{y}[n-1] - R_{yy}[1]|^2 \right\} &= \mathbb{E} \left\{ |y[n]|^2 |y[n-1]|^2 \right\} - |R_{yy}[1]|^2 \\ &= 2|A_0|^2 R_{vv}[0] + \mathbb{E} \left\{ |v[n]|^2 |v[n-1]|^2 \right\} - |R_{vv}[1]|^2 \\ &\quad + 2 \operatorname{Re} \left\{ \bar{A}_0^2 e^{-j\omega_0(2n-1)} \mathbb{E} \{ v[n]v[n-1] \} \right\} \\ &\quad + 2 \operatorname{Re} \left\{ \bar{A}_0 e^{-j\omega_0(n-1)} \mathbb{E} \left\{ |v[n]|^2 v[n-1] \right\} \right\} + 2 \operatorname{Re} \left\{ \bar{A}_0 e^{-j\omega_0 n} \mathbb{E} \{ v[n]|v[n-1]|^2 \} \right\} \end{aligned} \quad (\text{B.48})$$

and

$$\begin{aligned} \mathbb{E} \left\{ Q_c[n-1] (\bar{y}[n]y[n-1] - \bar{R}_{yy}[1]) \right\} &= \mathbb{E} \left\{ Q[n-1] \bar{y}[n]y[n-1] \right\} - |R_{yy}[1]|^2 \\ &= (1-\delta) \sum_{k=0}^{+\infty} \delta^k \mathbb{E} \left\{ y[n-1-k] \bar{y}[n-2-k] \bar{y}[n]y[n-1] \right\} - |R_{yy}[1]|^2 \\ &= (1-\delta) \sum_{k=0}^{+\infty} \delta^k \mathbb{E} \left\{ \bar{v}[n]v[n-1]v[n-1-k] \bar{v}[n-2-k] \right\} \\ &\quad + (1-\delta) \sum_{k=0}^{+\infty} \delta^k 2|A_0|^2 \operatorname{Re} \left\{ e^{-j\omega_0(k+1)} R_{vv}[k+1] \right\} - |R_{vv}[1]|^2 \\ &\quad + (1-\delta) \sum_{k=0}^{+\infty} \delta^k \bar{A}_0^2 e^{-j\omega_0(2n-2-k)} \mathbb{E} \{ v[n-1]v[n-1-k] \} \\ &\quad + (1-\delta) \sum_{k=0}^{+\infty} \delta^k A_0^2 e^{j\omega_0(2n-2-k)} \mathbb{E} \{ \bar{v}[n] \bar{v}[n-2-k] \} \\ &\quad + (1-\delta) \sum_{k=0}^{+\infty} \delta^k A_0 e^{j\omega_0(n-1)} \mathbb{E} \{ \bar{v}[n]v[n-1-k] \bar{v}[n-2-k] \} \\ &\quad + (1-\delta) \sum_{k=0}^{+\infty} \delta^k \bar{A}_0 e^{-j\omega_0(n-2-k)} \mathbb{E} \{ \bar{v}[n]v[n-1]v[n-1-k] \} \\ &\quad + (1-\delta) \sum_{k=0}^{+\infty} \delta^k A_0 e^{j\omega_0(n-1-k)} \mathbb{E} \{ \bar{v}[n]v[n-1] \bar{v}[n-2-k] \} \\ &\quad + (1-\delta) \sum_{k=0}^{+\infty} \delta^k \bar{A}_0 e^{-j\omega_0 n} \mathbb{E} \{ v[n-1]v[n-1-k] \bar{v}[n-2-k] \}, \end{aligned} \quad (\text{B.49})$$

where $Q[n]$ was expressed in terms of the infinite series $Q[n] = \sum_{k=0}^{+\infty} \delta^k y[n-k] \bar{v}[n-1-k]$.

From this point, two additional assumptions on the zero-mean white noise $v[n]$ are required to simplify the expression for the variance of $Q[n]$. First, $v[n]$ is supposed to follow a Gaussian distribution. In this case, all central moments of odd order are equal to zero [203]. Second, $v[n]$ is assumed to be a circular (or proper) complex random process [66]. Indeed, for circular processes, the pseudocorrelation, which is defined as

$$E\{v[n]v[n-k]\} = R_{v_R v_R}[k] - R_{v_I v_I}[k] + jR_{v_R v_I}[k] + jR_{v_I v_R}[k], \quad (\text{B.50})$$

with $v_R[n] = \text{Re}\{v[n]\}$ and $v_I[n] = \text{Im}\{v[n]\}$, is equal to zero. It is easy to show that the additive noises for the four considered input signals have indeed zero pseudocorrelations:

Noisy cisoid In this case, the noise $v_1[n]$ is a zero-mean complex circular Gaussian white process with variance σ^2 . Its pseudocorrelation is zero as the real and imaginary parts are independent and their variances are equal.

Noisy sinusoid For this input signal, the noise $v_2[n]$ is the analytic representation of a zero-mean real Gaussian white process $w_2[n]$,

$$v_2[n] = w_2[n] + jw_{2,h}[n] = w_2[n] + j(h[n] * w_2[n]),$$

where $w_{2,h}[n]$ is the DHT of $w_2[n]$, which follows a Gaussian distribution as well. Inserting (B.8) and (B.9) in (B.50), it is clear that $v_2[n]$ is a circular complex random process as its pseudocorrelation vanishes.

Filtered noisy cisoid Here, the noise $v[n]$ is obtained by filtering the circular Gaussian noise from the first case. Consequently, since the filter (B.2) is linear, the resulting process is also circular and Gaussian-distributed [66].

Filtered noisy sinusoid The same remarks as for the previous case apply to the filtered analytic representation of a zero-mean real Gaussian white noise.

Furthermore, since the noise $v[n]$ is assumed to follow a Gaussian distribution, its fourth order central moments can be computed using Isserlis' theorem [204]:

Theorem B.1 (Isserlis' theorem) *If $\xi_1, \xi_2, \dots, \xi_{2N+1}$ ($N = 1, 2, \dots$) are centered jointly Gaussian random variables (i.e. $E\{\xi_n\} = 0$ for every n), then*

$$E\{\xi_1 \xi_2 \cdots \xi_{2N}\} = \sum \prod E\{\xi_n \xi_m\} \quad (\text{B.51a})$$

and

$$E\{\xi_1 \xi_2 \cdots \xi_{2N+1}\} = 0, \quad (\text{B.51b})$$

where the notation $\sum \prod$ means summing over all distinct ways of partitioning $\xi_1, \xi_2, \dots, \xi_{2N}$ into pairs.

This theorem, which is also known as Wick's theorem [205], was first derived by Isserlis [206]. With this result which also holds for complex Gaussian random variables [207], the fourth order moment in (B.48) becomes

$$E\{|v[n]|^2 |v[n-1]|^2\} = R_{vv}^2[0] + |R_{vv}[1]|^2. \quad (\text{B.52})$$

and the one in (B.49) is now given by

$$E\{\bar{v}[n]v[n-1]v[n-1-k]\bar{v}[n-2-k]\} = |R_{vv}[1]|^2 + |R_{vv}[k+1]|^2. \quad (\text{B.53})$$

Finally, taking into account all the assumptions, the variance of $Q[n]$ is obtained as follows,

$$\begin{aligned} \text{Var}\{Q[n]\} = & \frac{1-\delta}{1+\delta} R_{vv}[0] \left(2|A_0|^2 + R_{vv}[0] \right) \\ & + 2\delta \frac{1-\delta}{1+\delta} \left(\sum_{k=0}^{+\infty} \delta^k |R_{vv}[k+1]|^2 + 2|A_0|^2 \sum_{k=0}^{+\infty} \delta^k \text{Re} \left\{ e^{-j\omega_0(k+1)} R_{vv}[k+1] \right\} \right). \end{aligned} \quad (\text{B.54})$$

The variances of $Q[n]$ and $\omega[n]$ were also investigated thoroughly with Monte Carlo simulations for the four considered input signals (noisy cisoid, noisy sinusoid, filtered noisy cisoid and filtered noisy sinusoid). This is particularly important for the instantaneous frequency estimate as no theoretical value could be calculated. The procedure for the simulations is identical to the one used for the bias (see Section B.3), and therefore only the most relevant details are repeated here. A simplified SFT without the adaptive filter (B.2) was applied to the first two input signals, while the complete algorithm was applied in the two other cases. The default parameters were chosen as follows: $A_0 = 1$, $\omega_0 = 0.4\pi$, $\text{SNR} = 10$ dB, $\delta = 0.95$, $\beta = 0.95$.

B.4.1 Noisy Cisoid

Inserting the autocorrelation of the noise $v_1[n]$ (B.4) into the expression for the variance of $Q[n]$ yields

$$\text{Var}\{Q[n]\} = \frac{1-\delta}{1+\delta} \sigma^2 (2|A_0|^2 + \sigma^2). \quad (\text{B.55})$$

This theoretical value was compared to the one obtained with Monte Carlo simulations. In addition, the variance of the instantaneous frequency estimate was also measured with numerical simulations. The influences of the SNR and the forgetting factor δ on the variances of $Q[n]$ and $\omega[n]$ were investigated as illustrated in Figure B.9. SNR values were taken from 0 to 20 dB in 1 dB steps (Figure B.9a). The Monte Carlo procedure confirmed that low noise variance leads to low estimation variance. The effect of the forgetting factor δ was analyzed by taking values between 0.5 and 0.95 in 0.05 steps (Figure B.9b). As expected, values of δ close to one also led to lower estimation variance. Furthermore, the theoretical and simulated variances of $Q[n]$ matched almost perfectly.

B.4.2 Noisy Sinusoid

Similarly to the previous case, the variance of $Q[n]$ for this input signal is obtained by combining the general expression (B.54) and the specific noise autocorrelation (B.9). The resulting variance is given by

$$\begin{aligned} \text{Var}\{Q[n]\} = & \frac{1-\delta}{1+\delta} 4\sigma^2 (|A_0|^2 + \sigma^2) \\ & + \frac{1-\delta}{1+\delta} \frac{8\sigma^2}{\pi} \left(\frac{2\sigma^2}{\pi} (\text{Li}_2(\delta) - \text{Li}_2(-\delta)) - |A_0|^2 \arg \left\{ \frac{1+\delta e^{-j\omega_0}}{1-\delta e^{-j\omega_0}} \right\} \right) \\ = & \frac{1-\delta}{1+\delta} 4\sigma^2 (|A_0|^2 + \sigma^2) + \frac{1-\delta}{1+\delta} \frac{8\sigma^2}{\pi} \left(\frac{4\sigma^2}{\pi} \chi_2(\delta) - |A_0|^2 \arg \left\{ \frac{1+\delta e^{-j\omega_0}}{1-\delta e^{-j\omega_0}} \right\} \right), \end{aligned} \quad (\text{B.56})$$

where $\text{Li}_s(r)$ is the polylogarithm (also known as Jonquière's function [208]) and $\chi_s(r)$ is Legendre chi function. These two functions are defined as [209]

$$\text{Li}_s(r) = \sum_{k=1}^{+\infty} \frac{r^k}{k^s} \quad (\text{B.57})$$

and

$$\chi_s(r) = \sum_{k=0}^{+\infty} \frac{r^{2k+1}}{(2k+1)^s}. \quad (\text{B.58})$$

This theoretical variance was then compared to the one measured with Monte Carlo simulations. The variance of $\omega[n]$ was estimated with the same procedure. Figure B.10 presents the outcomes of the simulations for various values of the parameters. The same remarks as for the first case regarding the influence of the noise level and the forgetting factor on the estimation variance apply (Figures B.10a and B.10b). Furthermore, as the sinusoid frequency is involved in (B.56), its effect on the variances of $Q[n]$ and $\omega[n]$ was also analyzed by varying ω_0 from 0.04π to 0.96π in 0.04π steps (Figure B.10c). One can observe that the frequency estimation variance is slightly smaller around 0.5π . It is reminiscent of the bias analysis where the bias magnitude was also smaller around this frequency (Figure B.6b). At last, theory and numerical simulations were in close agreement for this input signal as well.

B.4.3 Filtered Noisy Cisoid

In order to compute the variance of $Q[n]$ for this input signal, the autocorrelation of the noise (B.13) is inserted into (B.54):

$$\begin{aligned} \text{Var}\{Q[n]\} = & \frac{1-\delta}{1+\delta} \sigma^2 \frac{1-\beta}{1+\beta} \left(\frac{(1-\beta)^2}{1-2\beta \cos(\omega_0 - \omega_c) + \beta^2} \cdot 2|A_0|^2 + \sigma^2 \frac{1-\beta}{1+\beta} \right) \\ & + 2\delta \frac{1-\delta}{1+\delta} \sigma^2 \beta \frac{1-\beta}{1+\beta} \left(\frac{\sigma^2 \beta}{1-\delta\beta^2} \frac{1-\beta}{1+\beta} \right. \\ & \left. - \frac{(1-\beta)^2}{1-2\beta \cos(\omega_0 - \omega_c) + \beta^2} \cdot 2|A_0|^2 \frac{\cos(\omega_0 - \omega_c) - \delta\beta}{1-2\delta\beta \cos(\omega_0 - \omega_c) + \delta^2\beta^2} \right). \end{aligned} \quad (\text{B.59})$$

This variance and $\text{Var}\{\omega[n]\}$ were also computed with Monte Carlo simulations. They were measured with respect to the SNR, the forgetting factor, the central frequency and the bandwidth of the band-pass filter (B.2) as shown in Figure B.11. The noise level and δ had effects similar to the ones observed for the two input signals presented previously (Figures B.11a and B.11b). Like for the bias, the influence of the filter central frequency ω_c was investigated with a slightly modified SFT. Indeed, the central frequency was not adapted to the current frequency estimate, but instead it was set to fixed values taken from 0.04π to 0.96π in 0.04π steps. Clearly, the frequency estimation variance increased as the distance between the cisoid frequency and the central frequency increased (Figure B.11c). Thus, whenever the adaptive band-pass filter is not close enough to the cisoid frequency, the SFT provides an estimate with high bias and variance (Figure B.7b). Different values of the bandwidth parameter β also modified both $\text{Var}\{Q[n]\}$ and $\text{Var}\{\omega[n]\}$. This was assessed by varying β from 0.5 to 0.95 in 0.05 steps (Figure B.11d). As expected, values close to one led to narrower bandwidths and thus to lower estimation variances as more noise was filtered out. Overall, the theoretical and practical variances of $Q[n]$ matched closely in all these cases.

B.4.4 Filtered Noisy Sinusoid

Unlike the first three cases, the calculations for obtaining a closed form expression of the variance of $Q[n]$ for a filtered noisy sinusoid quickly become intractable due to the infinite series in (B.22) and (B.54). Nevertheless, a Monte Carlo procedure was applied for investigating the behavior of $\text{Var}\{Q[n]\}$ and $\text{Var}\{\omega[n]\}$ with respect to the different parameters. The results of the

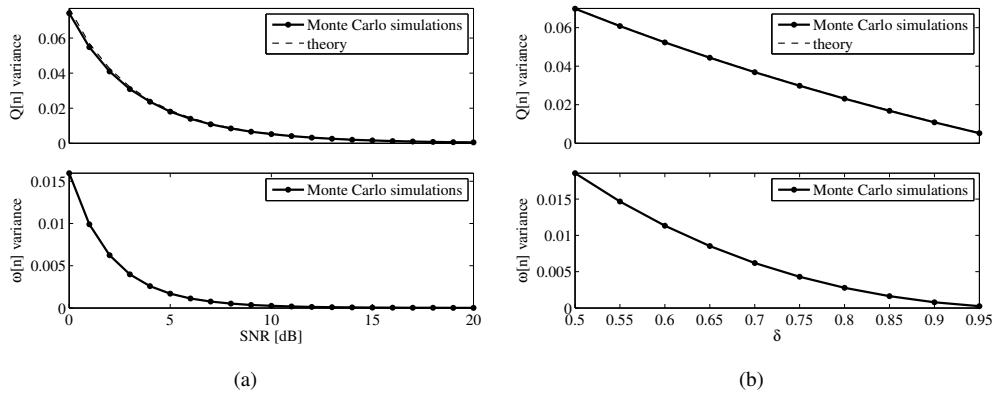


Figure B.9: Variance of the SFT for a noisy cisoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\delta = 0.95$) with respect (a) to the SNR, and (b) to the forgetting factor δ .

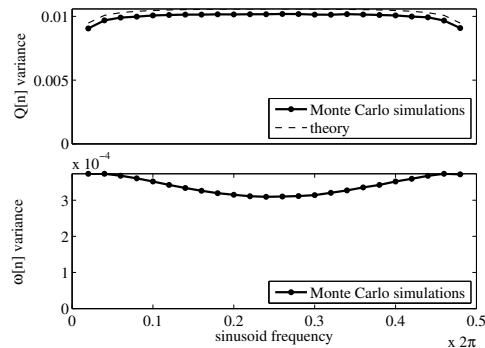
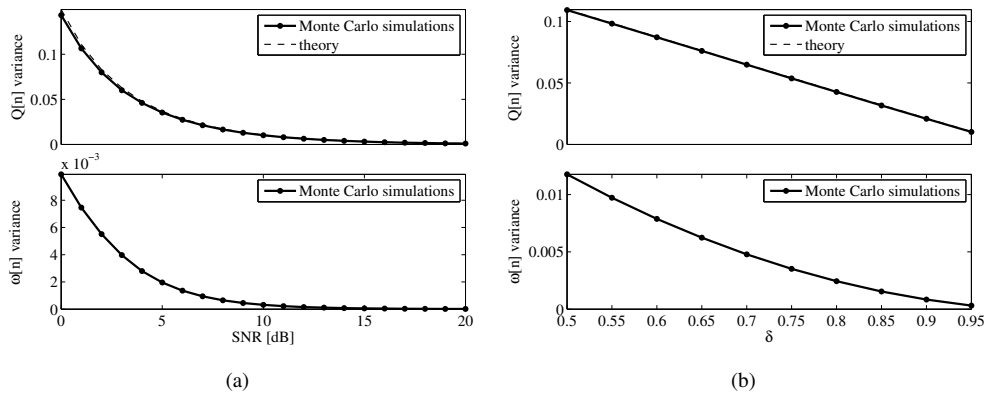


Figure B.10: Variance of the SFT for the analytic representation of a noisy sinusoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\delta = 0.95$) with respect (a) to the SNR, (b) to the forgetting factor δ , and (c) to the sinusoid frequency ω_0 .

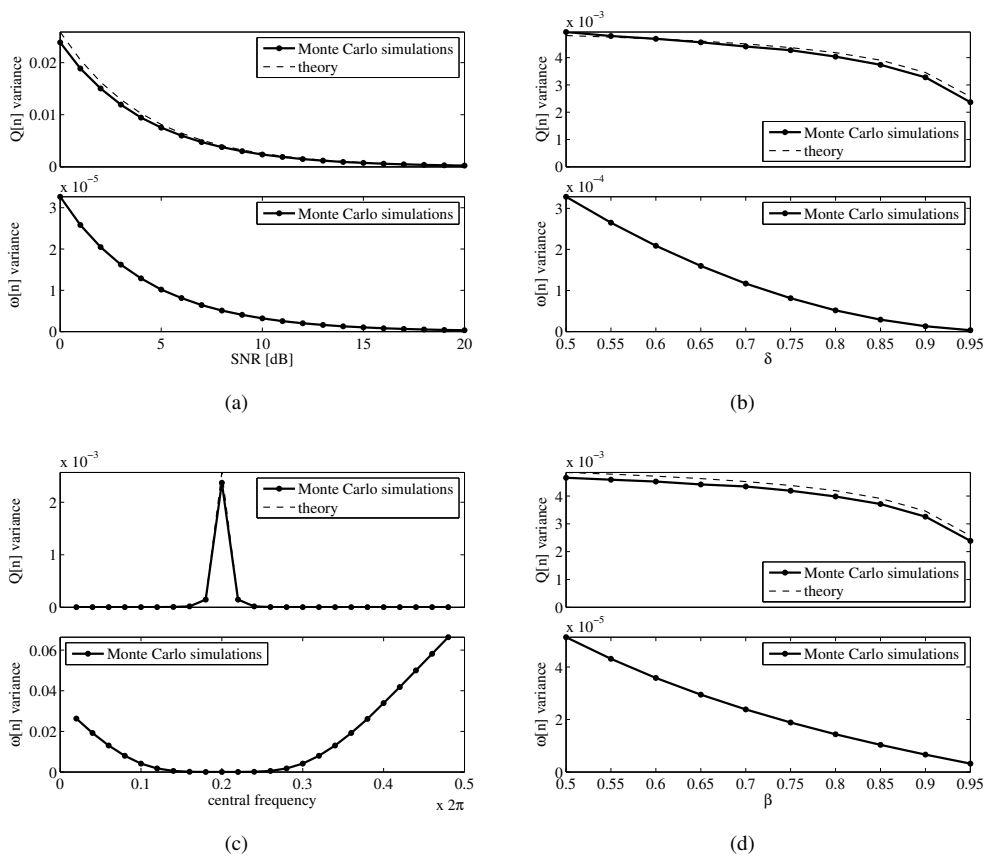


Figure B.11: Variance of the SFT for a filtered noisy cisoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$) with respect (a) to the SNR, (b) to the forgetting factor δ , (c) to the central frequency ω_c , and (d) to the bandwidth parameter β .

simulations are shown in Figure B.12 for varying SNRs, forgetting factors δ , sinusoid frequencies ω_0 , central frequencies ω_c and bandwidth parameters β of the band-pass filter. As for the other input signals, low noise variance led to low estimation variance (Figure B.12a). The forgetting factor and bandwidth parameter had the same effect as for the filtered noisy cisoid case: smaller variances were obtained for values close to one (Figures B.12b and B.12e). Furthermore, sinusoid frequencies close to 0 and π resulted in a slightly smaller variance for $\omega[n]$, but this effect was very limited (Figure B.12c). Finally, like for the previous input signal, the frequency estimation variance was minimized when the adaptive band-pass filter was centered on the true sinusoid frequency (Figure B.12d).

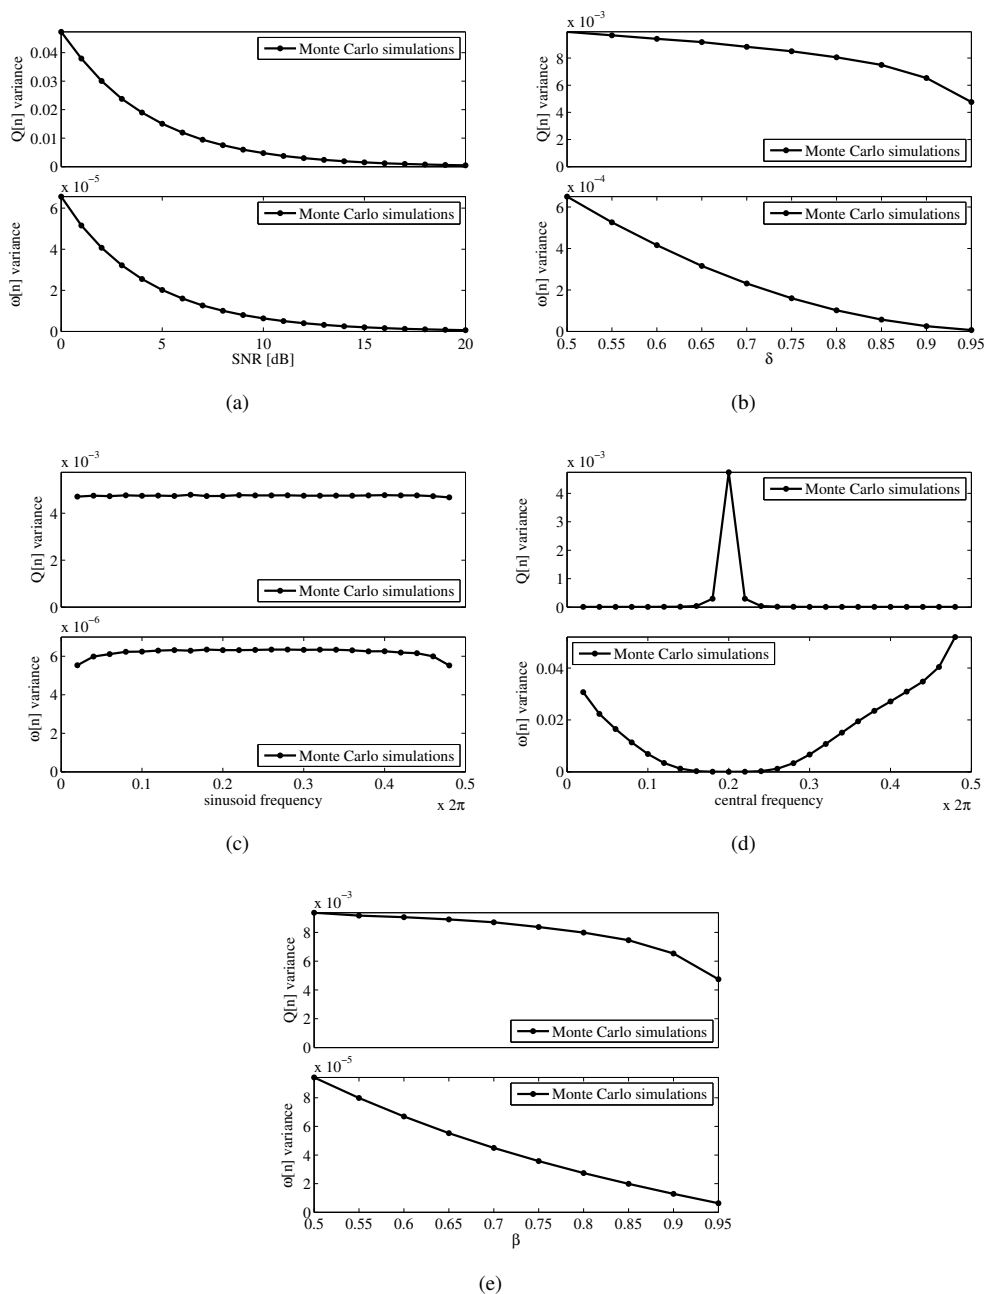


Figure B.12: Variance of the SFT for the filtered analytic representation of a noisy sinusoid ($A_0 = 1$, $\omega_0 = 0.4\pi$, SNR = 10 dB, $\beta = 0.95$, $\delta = 0.95$) with respect (a) to the SNR, (b) to the forgetting factor δ , (c) to the sinusoid frequency ω_0 , (d) to the central frequency ω_c , and (e) to the bandwidth parameter β .

Bibliography

- [1] P. Stoica and R. L. Moses, *Introduction to spectral analysis*. Prentice Hall, 1997.
- [2] P. D. Welch, “The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms,” *IEEE T Acoust Speech*, vol. AU-15, no. 2, pp. 70–73, June 1967.
- [3] M. H. Hayes, *Statistical digital signal processing and modeling*. John Wiley & Sons, 1996.
- [4] G. E. P. Box, G. M. Jenkins, and G. C. Reinsel, *Time series analysis: forecasting and control*, 4th ed. John Wiley & Sons, 2008.
- [5] V. F. Pisarenko, “The retrieval of harmonics from a covariance function,” *Geophys J Roy Astr S*, vol. 33, no. 3, pp. 347–366, September 1973.
- [6] R. O. Schmidt, “Multiple emitter location and signal parameter estimation,” *IEEE T Antenn Propag*, vol. AP-34, no. 3, pp. 276–280, March 1986.
- [7] R. Roy and T. Kailath, “ESPRIT — Estimation of signal parameters via rotational invariance techniques,” *IEEE T Acoust Speech*, vol. 37, no. 7, pp. 984–995, July 1989.
- [8] P. M. Clarkson, *Optimal and adaptive signal processing*. CRC Press, 1993.
- [9] M. Akay, Ed., *Time frequency and wavelets in biomedical signal processing*. IEEE Press, 1998.
- [10] P. Flandrin, *Time-frequency/time-scale analysis*. Academic Press, 1999.
- [11] B. Boashash, Ed., *Time frequency signal analysis and processing: a comprehensive reference*. Elsevier, 2003.
- [12] S. Mallat, *A wavelet tour of signal processing: the sparse way*. Academic Press, 2009.
- [13] M. Luise and R. Reggiannini, “Carrier frequency recovery in all-digital modems for burst-mode transmissions,” *IEEE T Commun*, vol. 43, no. 2/3/4, pp. 1169–1178, February/March/April 1995.
- [14] V. Lottici, M. Luise, C. Saccomando, and F. Spalla, “Blind carrier frequency tracking for filterbank multicarrier wireless communications,” *IEEE T Commun*, vol. 53, no. 10, pp. 1762–1772, October 2005.
- [15] H. Kawahara, I. Masuda-Katsuse, and A. de Cheveigné, “Restructuring speech representations using a pitch-adaptive time-frequency smoothing and an instantaneous-frequency-based F0 extraction: possible role of a repetitive structure in sounds,” *Speech Commun*, vol. 27, no. 3–4, pp. 187–207, April 1999.

- [16] A. Rao and R. Kumaresan, "On decomposing speech into modulated components," *IEEE T Speech Audi P*, vol. 8, no. 3, pp. 240–254, May 2000.
- [17] G. Evangelopoulos and P. Maragos, "Multiband modulation energy tracking for noisy speech detection," *IEEE T Acoust Speech*, vol. 14, no. 6, pp. 2024–2038, November 2006.
- [18] M. Ferdjallah and R. E. Barr, "Adaptive digital notch filter design on the unit circle for the removal of powerline noise from biomedical signals," *IEEE T Bio-Med Eng*, vol. 41, no. 6, pp. 529–536, June 1994.
- [19] D. P. Nguyen, M. A. Wilson, E. N. Brown, and R. Barbieri, "Measuring instantaneous frequency of local field potential oscillations using the Kalman smoother," *J Neurosci Meth*, vol. 184, no. 2, pp. 365–374, November 2009.
- [20] B. Boashash, "Estimating and interpreting the instantaneous frequency of a signal — Part 1: fundamentals," *P IEEE*, vol. 80, no. 4, pp. 520–538, April 1992.
- [21] J. R. Carson and T. C. Fry, "Variable frequency electric circuit theory with application to the theory of frequency-modulation," *Bell Syst Tech J*, vol. 16, no. 4, pp. 513–540, October 1937.
- [22] B. van der Pol, "The fundamental principles of frequency modulation," *J IEE*, vol. 93, no. 23, pp. 153–158, May 1946.
- [23] D. Gabor, "Theory of communication," *J IEE*, vol. 93, no. 26, pp. 429–457, November 1946.
- [24] D. W. Kammler, *A first course in Fourier analysis*. Prentice Hall, 2000.
- [25] J. Ville, "Théorie et applications de la notion de signal analytique," *Câbles et transmissions*, vol. 2, no. 1, pp. 61–74, January 1948.
- [26] J. Shekel, "Instantaneous frequency," *P IRE*, vol. 41, no. 4, p. 548, April 1953.
- [27] M. Chavez, M. Besserve, C. Adam, and J. Martinerie, "Towards a proper estimation of phase synchronization from time series," *J Neurosci Meth*, vol. 154, no. 1–2, pp. 149–160, June 2006.
- [28] L. Mandel, "Interpretation of instantaneous frequencies," *Am J Phys*, vol. 42, no. 10, pp. 840–846, October 1974.
- [29] P. J. Loughlin and B. Tacer, "Comments on the interpretation of instantaneous frequency," *IEEE Signal Proc Let*, vol. 4, no. 5, pp. 123–125, May 1997.
- [30] W. Nho and P. J. Loughlin, "When is the instantaneous frequency the average frequency at each time?" *IEEE Signal Proc Let*, vol. 6, no. 4, pp. 78–80, April 1999.
- [31] A. V. Oppenheim and R. W. Schaffer, *Discrete-time signal processing*, 3rd ed. Pearson, 2010.
- [32] L. Grafakos, "An elementary proof of the square summability of the discrete Hilbert transform," *Am Math Mon*, vol. 101, no. 5, pp. 456–458, May 1994.
- [33] V. Čížek, "Discrete Hilbert transform," *IEEE T Acoust Speech*, vol. 18, no. 4, pp. 340–343, December 1970.
- [34] S. L. Marple, "Computing the discrete-time analytic signal from the FFT," *IEEE T Signal Proces*, vol. 47, no. 9, pp. 2600–2603, September 1999.

- [35] B. Boashash, "Estimating and interpreting the instantaneous frequency of a signal — Part 2: algorithms and applications," *P IEEE*, vol. 80, no. 4, pp. 540–568, April 1992.
- [36] H. M. Teager, "Some observations on oral air flow during phonation," *IEEE T Acoust Speech*, vol. ASSP-28, no. 5, pp. 599–601, October 1980.
- [37] H. M. Teager and S. M. Teager, "A phenomenological model for vowel production in the vocal tract," in *Speech sciences: recent advances*, R. G. Daniloff, Ed. College-Hill Press, 1985, pp. 73–109.
- [38] —, "Evidence for nonlinear sound production mechanisms in the vocal tract," in *Speech production and speech modelling*, W. J. Hardcastle and A. Marchal, Eds. Kluwer Academic Publishers, 1990, pp. 241–261.
- [39] J. F. Kaiser, "On a simple algorithm to calculate the 'energy' of a signal," in *Int Conf Acoust Spee*, vol. 1, April 1990, pp. 381–384.
- [40] P. Maragos, J. F. Kaiser, and T. F. Quatieri, "On amplitude and frequency demodulation using energy operators," *IEEE T Signal Proces*, vol. 41, no. 4, pp. 1532–1550, April 1993.
- [41] —, "Energy separation in signal modulations with application to speech processing," *IEEE T Signal Proces*, vol. 41, no. 10, pp. 3024–3051, October 1993.
- [42] B. Widrow, J. R. Glover, Jr., J. M. McCool, J. Kaunitz, C. S. Williams, R. H. Hearn, J. R. Zeidler, E. Dong, Jr., and R. C. Goodlin, "Adaptive noise cancelling: principles and applications," *P IEEE*, vol. 63, no. 12, pp. 1692–1716, December 1975.
- [43] J. R. Zeidler, E. H. Satorius, D. M. Chabries, and H. T. Wexler, "Adaptive enhancement of multiple sinusoids in uncorrelated noise," *IEEE T Acoust Speech*, vol. ASSP-26, no. 3, pp. 240–254, June 1978.
- [44] J. R. Treichler, "Transient and convergent behavior of the adaptive line enhancer," *IEEE T Acoust Speech*, vol. ASSP-27, no. 1, pp. 53–62, February 1979.
- [45] N. J. Bershad and P. L. Feintuch, "The recursive adaptive LMS filter — A line enhancer application and analytical model for the mean weight behavior," *IEEE T Acoust Speech*, vol. ASSP-28, no. 6, pp. 652–660, December 1980.
- [46] V. U. Reddy, B. Egardt, and T. Kailath, "Optimized lattice-form adaptive line enhancer for a sinusoidal signal in broad-band noise," *IEEE T Acoust Speech*, vol. ASSP-29, no. 3, pp. 702–710, June 1981.
- [47] B. Friedlander, "A recursive maximum likelihood algorithm for ARMA line enhancement," *IEEE T Acoust Speech*, vol. ASSP-30, no. 4, pp. 651–657, August 1982.
- [48] A. Nehorai, "A minimal parameter adaptive notch filter with constrained poles and zeros," *IEEE T Acoust Speech*, vol. 33, no. 4, pp. 983–996, August 1985.
- [49] D. R. Hush, N. Ahmed, R. David, and S. D. Stearns, "An adaptive IIR structure for sinusoidal enhancement, frequency estimation, and detection," *IEEE T Acoust Speech*, vol. ASSP-34, no. 6, pp. 1380–1390, December 1986.
- [50] D. P. Bertsekas, *Nonlinear programming*, 2nd ed. Athena Scientific, 1999.
- [51] H.-E. Liao, "Two discrete oscillator based adaptive notch filters (OSC ANFs) for noisy sinusoids," *IEEE T Signal Proces*, vol. 53, no. 2, pp. 528–538, February 2005.

- [52] N. I. Cho and S. U. Lee, "On the adaptive lattice notch filter for the detection of sinusoids," *IEEE T Circuits-II*, vol. 40, no. 7, pp. 405–416, July 1993.
- [53] M. Niedźwiecki and P. Kaczmarek, "Tracking analysis of an generalized adaptive notch filter," *IEEE T Signal Proces*, vol. 54, no. 1, pp. 304–314, January 2006.
- [54] —, "Generalized adaptive notch filter with a self-optimization capability," *IEEE T Signal Proces*, vol. 54, no. 11, pp. 4185–4193, November 2006.
- [55] P. Tichavský, "Posterior Cramér-Rao bound for adaptive harmonic retrieval," *IEEE T Signal Proces*, vol. 43, no. 5, pp. 1299–1302, May 1995.
- [56] P. A. Regalia, "A complex adaptive notch filter," *IEEE Signal Proc Let*, vol. 17, no. 11, pp. 937–940, November 2010.
- [57] P. P. Vaidyanathan, P. A. Regalia, and S. K. Mitra, "Design of doubly-complementary IIR digital filters using a single complex allpass filter, with multirate applications," *IEEE T Circuits Syst*, vol. 34, no. 4, pp. 378–389, April 1987.
- [58] H. C. So, "A comparative study of three recursive least-squares algorithms for single-tone frequency tracking," *Signal Process*, vol. 83, no. 9, pp. 2059–2062, September 2003.
- [59] H. C. So, K. W. Chan, Y. T. Chan, and K. C. Ho, "Linear prediction approach for efficient frequency estimation of multiple real sinusoids: algorithms and analyses," *IEEE T Signal Proces*, vol. 53, no. 7, pp. 2290–2305, July 2005.
- [60] J. Yang, H. Xi, and W. Guo, "Robust modified Newton algorithm for adaptive frequency estimation," *IEEE Signal Proc Let*, vol. 14, no. 11, pp. 879–882, November 2007.
- [61] Y. T. Chan, J. M. M. Lavoie, and J. B. Plant, "A parameter estimation approach to estimation of frequencies of sinusoids," *IEEE T Acoust Speech*, vol. ASSP-29, no. 2, pp. 214–219, April 1981.
- [62] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge University Press, 1985.
- [63] L. Uldry, C. Duchêne, Y. Prudat, M. M. Murray, and J.-M. Vesin, "Adaptive tracking of EEG frequency components," in *Advanced biosignal processing*, A. Nait-Ali, Ed. Springer, 2009, pp. 123–144.
- [64] J. Van Zaen, L. Uldry, C. Duchêne, Y. Prudat, R. A. Meuli, M. M. Murray, and J.-M. Vesin, "Adaptive tracking of EEG oscillations," *J Neurosci Meth*, vol. 186, no. 1, pp. 97–106, January 2010.
- [65] S. Haykin, *Adaptive filter theory*, 4th ed. Prentice Hall, 2001.
- [66] F. D. Neeser and J. L. Massey, "Proper complex random processes with applications to information theory," *IEEE T Inform Theory*, vol. 39, no. 4, pp. 1293–1302, July 1993.
- [67] J. Van Zaen, L. Uldry, A. Buttu, L. Kappenberger, and J.-M. Vesin, "Harmonic frequency tracking algorithm for predicting the success of pharmacological cardioversion of atrial fibrillation," *Europace*, vol. 13, no. suppl 3, p. P513, June 2011.
- [68] A. Buttu, J. Van Zaen, A. Viso, A. Forclaz, P. Pascale, S. M. Narayan, J.-M. Vesin, and E. Pruvot, "A new method for ECG tracking of persistent atrial fibrillation termination during stepwise ablation," in *Comput Cardiol*, vol. 38, September 2011, pp. 165–168.

- [69] A. Buttu, J. Van Zaen, A. Viso, A. Forclaz, P. Pascale, S. M. Narayan, E. Pruvot, and J.-M. Vesin, "Adaptive frequency tracking of the baseline ECG identifies the site of atrial fibrillation termination by catheter ablation," *Biomed Signal Proces*, 2012, submitted.
- [70] G. Buzsáki, *Rhythms of the brain*. Oxford University Press, 2006.
- [71] J. R. Cameron, W. T. Thomson, and A. B. Dow, "Vibration and current monitoring for detecting airgap eccentricity in large induction motors," *IEE Proc-B*, vol. 133, no. 3, pp. 155–163, May 1986.
- [72] M. El Hachemi Benbouzid, "A review of induction motors signature analysis as a medium for faults detection," *IEEE T Ind Electron*, vol. 47, no. 5, pp. 984–993, October 2000.
- [73] Y. Prudat and J.-M. Vesin, "Multi-signal extension of adaptive frequency tracking algorithms," *Signal Process*, vol. 89, no. 6, pp. 963–973, June 2009.
- [74] Y. Prudat, "Adaptive frequency tracking and application to biomedical signals," Ph.D. dissertation, École polytechnique fédérale de Lausanne, 2009, director: J.-M. Vesin.
- [75] S. Roman, *Introduction to the mathematics of finance : from risk management to options pricing*. Springer, 2004.
- [76] B. Vidakovic, *Statistics for bioengineering sciences*. Springer, 2011.
- [77] J. L. Flanagan and R. M. Golden, "Phase vocoder," *Bell Syst Tech J*, vol. 45, no. 9, pp. 1493–1509, November 1966.
- [78] P. Tichavský and A. Nehorai, "Comparative study of four adaptive frequency trackers," *IEEE T Signal Proces*, vol. 45, no. 6, pp. 1473–1484, June 1997.
- [79] S. Kay, "Statistically/computationally efficient frequency estimation," in *Int Conf Acoust Spee*, vol. 4, April 1988, pp. 2292–2295.
- [80] M. Niedźwiecki and A. Sobociński, "A simple way of increasing estimation accuracy of generalized adaptive notch filters," *IEEE Signal Proc Let*, vol. 14, no. 3, pp. 217–220, March 2007.
- [81] M. Niedźwiecki, "Statistically efficient smoothing algorithm for time-varying frequency estimation," *IEEE T Signal Proces*, vol. 56, no. 8, pp. 3846–3854, August 2008.
- [82] —, "Generalized adaptive notch smoothing revisited," *IEEE T Signal Proces*, vol. 58, no. 3, pp. 1565–1576, March 2010.
- [83] M. Niedźwiecki and M. Meller, "New algorithms for adaptive notch smoothing," *IEEE T Signal Proces*, vol. 59, no. 5, pp. 2024–2037, May 2011.
- [84] F. J. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *P IEEE*, vol. 66, no. 1, pp. 51–83, January 1978.
- [85] M. M. Murray, G. R. Wylie, B. A. Higgins, D. C. Javitt, C. E. Schroeder, and J. J. Foxe, "The spatiotemporal dynamics of illusory contour processing: combined high-density electrical mapping, source analysis, and functional magnetic resonance imaging," *J Neurosci*, vol. 22, no. 12, pp. 5055–5073, June 2002.

- [86] D. Laude, J.-L. Elghozi, A. Girard, E. Bellard, M. Bouhaddi, P. Castiglioni, C. Cerutti, A. Cividjian, M. Di Rienzo, J.-O. Fortrat, B. Janssen, J. M. Karemaker, G. Lefthériotis, G. Parati, P. B. Persson, A. Porta, L. Quintin, J. Regnard, H. Rüdiger, and H. M. Stauss, "Comparison of various techniques used to estimate spontaneous baroreflex sensitivity (the EuroBaVar study)," *Am J Physiol-Reg I*, vol. 286, no. 1, pp. R226–R231, January 2004.
- [87] M. Malik, J. T. Bigger, A. J. Camm, R. E. Kleiger, A. Malliani, A. J. Moss, and P. J. Schwartz, "Heart rate variability," *Eur Heart J*, vol. 17, no. 3, pp. 354–381, March 1996.
- [88] H. Berger, "Über das Elektrenkephalogramm des Menschen," *Eur Arch Psy Clin N*, vol. 87, no. 1, pp. 527–570, December 1929.
- [89] P. L. Nunez and R. Srinivasan, *Electric fields of the brain: the neurophysics of EEG*, 2nd ed. Oxford University Press, 2006.
- [90] E. Niedermeyer and F. Lopes Da Silva, Eds., *Electroencephalography: basic principles, clinical applications, and related fields*, 5th ed. Lippincott Williams & Wilkins, 2005.
- [91] D. Brandeis, C. M. Michel, and F. Amzica, "From neuronal activity to scalp potential fields," in *Electrical neuroimaging*, C. M. Michel, T. Koenig, D. Brandeis, L. R. R. Gianotti, and J. Wackermann, Eds. Cambridge University Press, 2009, pp. 1–24.
- [92] P. Achermann and A. A. Borbély, "Low-frequency (<1 Hz) oscillations in the human sleep electroencephalogram," *Neuroscience*, vol. 81, no. 1, pp. 213–222, August 1997.
- [93] E. Kirmizi-Alsan, Z. Bayraktaroglu, H. Gurvit, Y. H. Keskin, M. Emre, and T. Demiralp, "Comparative analysis of event-related potentials during Go/NoGo and CPT: decomposition of electrophysiological markers of response inhibition and sustained attention," *Brain Res*, vol. 1104, no. 1, pp. 114–128, August 2006.
- [94] A. D. Ekstrom, J. B. Caplan, E. Ho, K. Shattuck, I. Fried, and M. J. Kahana, "Human hippocampal theta activity during virtual navigation," *Hippocampus*, vol. 15, no. 7, pp. 881–889, 2005.
- [95] O. Jensen and C. D. Tesche, "Frontal theta activity in humans increases with memory load in a working memory task," *Eur J Neurosci*, vol. 15, no. 8, pp. 1395–1399, April 2002.
- [96] W. Klimesch, S. Hanslmayr, P. Sauseng, W. Gruber, C. J. Brozinsky, N. E. A. Kroll, A. P. Yonelinas, and M. Doppelmayr, "Oscillatory EEG correlates of episodic trace decay," *Cereb Cortex*, vol. 16, no. 2, pp. 280–290, February 2006.
- [97] W. Klimesch, P. Sauseng, and S. Hanslmayr, "EEG alpha oscillations: the inhibition–timing hypothesis," *Brain Res Rev*, vol. 53, no. 1, pp. 63–88, January 2007.
- [98] S. Palva and J. M. Palva, "New vistas for α -frequency band oscillations," *Trends Neurosci*, vol. 30, no. 4, pp. 150–158, April 2007.
- [99] S. Salenius and R. Hari, "Synchronous cortical oscillatory activity during motor action," *Curr Opin Neurobiol*, vol. 13, no. 6, pp. 678–684, December 2003.
- [100] C. Tallon-Baudry, O. Bertrand, and C. Fischer, "Oscillatory synchrony between human extrastriate areas during visual short-term memory maintenance," *J Neurosci*, vol. 21, no. 20, p. RC177, October 2001.

- [101] M.-P. Deiber, P. Missonnier, O. Bertrand, G. Gold, L. Fazio-Costa, V. Ibañez, and P. Gianakopoulos, "Distinction between perceptual and attentional processing in working memory tasks: a study of phase-locked and induced oscillatory brain dynamics," *J Cognitive Neurosci*, vol. 19, no. 1, pp. 158–172, January 2007.
- [102] T. Womelsdorf, P. Fries, P. P. Mitra, and D. Robert, "Gamma-band synchronization in visual cortex predicts speed of change detection," *Nature*, vol. 439, no. 7077, pp. 733–736, February 2006.
- [103] E. Edwards, M. Soltani, L. Y. Deouell, M. S. Berger, and R. T. Knight, "High gamma activity in response to deviant auditory stimuli recorded directly from human cortex," *J Neurophysiol*, vol. 94, no. 6, pp. 4269–4280, December 2005.
- [104] J.-M. Schoffelen, R. Oostenveld, and P. Fries, "Neuronal coherence as a mechanism of effective corticospinal interaction," *Science*, vol. 308, no. 5718, pp. 111–113, April 2005.
- [105] C. S. Herrmann, M. H. J. Munk, and A. K. Engel, "Cognitive functions of gamma-band activity: memory match and utilization," *Trends Cogn Sci*, vol. 8, no. 8, pp. 347–355, August 2004.
- [106] S. L. Gonzalez Andino, C. M. Michel, G. Thut, T. Landis, and R. Grave de Peralta, "Prediction of response speed by anticipatory high-frequency (gamma band) oscillations in the human brain," *Hum Brain Mapp*, vol. 24, no. 1, pp. 50–58, January 2005.
- [107] M. Bauer, R. Oostenveld, M. Peeters, and P. Fries, "Tactile spatial attention enhances gamma-band activity in somatosensory cortex and reduces low-frequency activity in parieto-occipital areas," *JNeurosci*, vol. 26, no. 2, pp. 490–501, January 2006.
- [108] J. Kaiser and W. Lutzenberger, "Human gamma-band activity: a window to cognitive processing," *Neuroreport*, vol. 16, no. 3, pp. 207–211, February 2005.
- [109] O. Jensen and L. L. Colgin, "Cross-frequency coupling between neuronal oscillations," *Trends Cogn Sci*, vol. 11, no. 7, pp. 267–269, July 2007.
- [110] J. Van Zaen, M. M. Murray, R. A. Meuli, and J.-M. Vesin, "Adaptive filtering methods for identifying cross-frequency couplings in human EEG," *PLoS One*, 2012, submitted.
- [111] F. Varela, J.-P. Lachaux, E. Rodriguez, and J. Martinerie, "The brainweb: phase synchronization and large-scale integration," *Nat Rev Neurosci*, vol. 2, no. 4, pp. 229–239, April 2001.
- [112] A. K. Engel, P. Fries, and W. Singer, "Dynamic predictions: oscillations and synchrony in top-down processing," *Nat Rev Neurosci*, vol. 2, no. 10, pp. 704–716, October 2001.
- [113] J. Fell and N. Axmacher, "The role of phase synchronization in memory processes," *Nat Rev Neurosci*, vol. 12, no. 2, pp. 105–118, February 2011.
- [114] P. Fries, "A mechanism for cognitive dynamics: neuronal communication through neuronal coherence," *Trends Cogn Sci*, vol. 9, no. 10, pp. 474–480, October 2005.
- [115] N. Axmacher, F. Mormann, G. Fernández, C. E. Elger, and J. Fell, "Memory formation by neuronal synchronization," *Brain Res Rev*, vol. 52, no. 1, pp. 170–182, August 2006.
- [116] W. Singer and C. M. Gray, "Visual feature integration and the temporal correlation hypothesis," *Annu Rev Neurosci*, vol. 18, pp. 555–586, 1995.

- [117] A. K. Engel and W. Singer, "Temporal binding and the neural correlates of sensory awareness," *Trends Cogn Sci*, vol. 5, no. 1, pp. 16–25, January 2001.
- [118] A. von Stein and J. Sarnthein, "Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization," *Int J Psychophysiol*, vol. 38, no. 3, pp. 301–313, December 2000.
- [119] R. T. Canolty, E. Edwards, S. S. Dalal, M. Soltani, S. S. Nagarajan, H. E. Kirsch, M. S. Berger, N. M. Barbaro, and R. T. Knight, "High gamma power is phase-locked to theta oscillations in human neocortex," *Science*, vol. 313, no. 5793, pp. 1626–1628, September 2006.
- [120] R. T. Canolty, K. Ganguly, S. W. Kennerley, C. F. Cadieu, K. Koepsell, J. D. Wallis, and J. M. Carmena, "Oscillatory phase coupling coordinates anatomically dispersed functional cell assemblies," *P Natl Acad Sci USA*, vol. 107, no. 40, pp. 17 356–17 361, October 2010.
- [121] P. Lakatos, C.-M. Chen, M. N. O'Connell, A. Mills, and C. E. Schroeder, "Neuronal oscillations and multisensory interaction in primary auditory cortex," *Neuron*, vol. 53, no. 2, pp. 279–292, January 2007.
- [122] P. Lakatos, G. Karmos, A. D. Mehta, I. Ulbert, and C. E. Schroeder, "Entrainment of neuronal oscillations as a mechanism of attentional selection," *Science*, vol. 320, no. 5872, pp. 110–113, April 2008.
- [123] T. Demiralp, Z. Bayraktaroglu, D. Lenz, S. Junge, N. A. Busch, B. Maess, M. Ergen, and C. S. Herrmann, "Gamma amplitudes are coupled to theta phase in human EEG during visual perception," *Int J Psychophysiol*, vol. 64, no. 1, pp. 24–30, April 2007.
- [124] P. Sauseng, W. Klimesch, K. F. Heise, W. R. Gruber, E. Holz, A. A. Karim, M. Glennon, C. Gerloff, N. Birbaumer, and F. C. Hummel, "Brain oscillatory substrates of visual short-term memory capacity," *Curr Biol*, vol. 19, no. 21, pp. 1846–1852, November 2009.
- [125] B. L. Foster and J. Parvizi, "Resting oscillations and cross-frequency coupling in the human posteromedial cortex," *Neuroimage*, vol. 60, no. 1, pp. 384–391, March 2012.
- [126] G. Buzsáki and X.-J. Wang, "Mechanisms of gamma oscillations," *Annu Rev Neurosci*, vol. 35, no. 1, pp. 203–225, July 2012.
- [127] G. Kanizsa, "Subjective contours," *Sci Am*, vol. 234, no. 4, pp. 48–52, April 1976.
- [128] C. Tallon-Baudry, O. Bertrand, C. Delpuech, and J. Pernier, "Stimulus specificity of phase-locked and non-phase-locked 40 Hz visual responses in human," *J Neurosci*, vol. 16, no. 13, pp. 4240–4249, July 1996.
- [129] C. S. Herrmann, A. Mecklinger, and E. Pfeifer, "Gamma responses and ERPs in a visual classification task," *Clin Neurophysiol*, vol. 110, no. 4, pp. 636–642, April 1999.
- [130] G. Csibra, G. Davis, M. W. Spratling, and M. H. Johnson, "Gamma oscillations and object processing in the infant brain," *Science*, vol. 290, no. 5496, pp. 1582–1585, November 2000.
- [131] M. M. Murray, D. M. Foxe, D. C. Javitt, and J. J. Foxe, "Setting boundaries: brain dynamics of modal and amodal illusory shape completion in humans," *J Neurosci*, vol. 24, no. 31, pp. 6898–6903, August 2004.

- [132] M. M. Murray, M. L. Imber, D. C. Javitt, and J. J. Foxe, "Boundary completion is automatic and dissociable from shape discrimination," *J Neurosci*, vol. 26, no. 46, pp. 12 043–12 054, November 2006.
- [133] M. Shpaner, M. M. Murray, and J. J. Foxe, "Early processing in the human lateral occipital complex is highly responsive to illusory contours but not to salient regions," *Eur J Neurosci*, vol. 30, no. 10, pp. 2018–2028, November 2009.
- [134] J.-F. Knebel and M. M. Murray, "Towards a resolution of conflicting models of illusory contour processing in humans," *Neuroimage*, vol. 59, no. 3, pp. 2808–2817, February 2012.
- [135] C. A. Bosman, F. Zamorano, and F. Aboitiz, "Functional differences of low- and high-frequency oscillatory dynamics during illusory border perception," *Brain Res*, vol. 1319, pp. 92–102, March 2010.
- [136] R. Quian Quiroga, A. Kraskov, T. Kreuz, and P. Grassberger, "Performance of different synchronization measures in real data: a case study on electroencephalographic signals," *Phys Rev E*, vol. 65, no. 4, p. 041903, March 2002.
- [137] P. Tass, M. G. Rosenblum, J. Weule, J. Kurths, A. Pikovsky, J. Volkmann, A. Schnitzler, and H.-J. Freund, "Detection of $n : m$ phase locking from noisy data: application to magnetoencephalography," *Phys Rev Lett*, vol. 81, no. 15, pp. 3291–3294, October 1998.
- [138] J.-P. Lachaux, E. Rodriguez, J. Martinerie, and F. J. Varela, "Measuring phase synchrony in brain signals," *Hum Brain Mapp*, vol. 8, no. 4, pp. 194–208, 1999.
- [139] P. G. Schyns, G. Thut, and J. Gross, "Cracking the code of oscillatory activity," *PLoS Biol*, vol. 9, no. 5, p. e1001064, May 2011.
- [140] P. Celka, "Statistical analysis of the phase-locking value," *IEEE Signal Proc Let*, vol. 14, no. 9, pp. 577–580, September 2007.
- [141] R. C. Oldfield, "The assessment and analysis of handedness: the Edinburgh inventory," *Neuropsychologia*, vol. 9, no. 1, pp. 97–113, March 1971.
- [142] D. Brunet, M. M. Murray, and C. M. Michel, "Spatiotemporal Analysis of Multichannel EEG: CARTOOL," *Comput Intell Neurosci*, vol. 2011, p. 813870, 2011.
- [143] T. Schreiber and A. Schmitz, "Surrogate time series," *Physica D*, vol. 142, no. 3–4, pp. 346–382, August 2000.
- [144] P. Borgnat, P. Flandrin, P. Honeine, C. Richard, and J. Xiao, "Testing stationarity with surrogates: a time-frequency approach," *IEEE T Signal Proces*, vol. 58, no. 7, pp. 3459–3470, July 2010.
- [145] R. R. Hocking, *Methods and applications of linear models: regression and the analysis of variance*, 2nd ed. John Wiley & Sons, 2003.
- [146] S. M. Ross, *Introductory statistics*, 3rd ed. Academic Press, 2010.
- [147] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*. Chapman & Hall, 1993.
- [148] P. Billingsley, *Probability and measure*, 3rd ed. John Wiley & Sons, 1995.
- [149] W. D. Penny, E. Duzel, K. J. Miller, and J. G. Ojemann, "Testing for nested oscillation," *J Neurosci Meth*, vol. 174, no. 1, pp. 50–61, September 2008.

- [150] F. P. de Lange, O. Jensen, M. Bauer, and I. Toni, "Interactions between posterior gamma and frontal alpha/beta oscillations during imagined actions," *Front Hum Neurosci*, vol. 2, p. 7, August 2008.
- [151] D. Osipova, D. Hermes, and O. Jensen, "Gamma power is phase-locked to posterior alpha activity," *PLoS ONE*, vol. 3, no. 12, p. e3990, December 2008.
- [152] C. E. Schroeder and P. Lakatos, "Low-frequency neuronal oscillations as instruments of sensory selection," *Trends Neurosci*, vol. 32, no. 1, pp. 9–18, January 2009.
- [153] J. J. Foxe, M. M. Murray, and D. C. Javitt, "Filling-in in schizophrenia: a high-density electrical mapping and source-analysis investigation of illusory contour processing," *Cereb Cortex*, vol. 15, no. 12, pp. 1914–1927, December 2005.
- [154] J.-F. Knebel, D. C. Javitt, and M. M. Murray, "Impaired early visual response modulations to spatial information in chronic schizophrenia," *Psychiat Res-Neuroim*, vol. 193, no. 3, pp. 168–176, September 2011.
- [155] C. Tallon-Baudry and O. Bertrand, "Oscillatory gamma activity in humans and its role in object representation," *Trends Cogn Sci*, vol. 3, no. 4, pp. 151–162, April 1999.
- [156] G. Buzsáki and A. Draguhn, "Neuronal oscillations in cortical networks," *Science*, vol. 304, no. 5679, pp. 1926–1929, June 2004.
- [157] R. T. Canolty and R. T. Knight, "The functional role of cross-frequency coupling," *Trends Cogn Sci*, vol. 14, no. 11, pp. 506–515, November 2010.
- [158] C. M. Michel and M. M. Murray, "Towards the utilization of EEG as a brain imaging tool," *Neuroimage*, vol. 61, no. 2, pp. 371–385, June 2012.
- [159] L. Astolfi, J. Toppi, F. De Vico Fallani, G. Vecchiato, S. Salinari, D. Mattia, F. Cincotti, and F. Babiloni, "Neuroelectrical hyperscanning measures simultaneous brain activity in humans," *Brain Topogr*, vol. 23, no. 3, pp. 243–256, September 2010.
- [160] F. De Vico Fallani, F. Rodrigues, L. da Fontoura Costa, L. Astolfi, F. Cincotti, D. Mattia, S. Salinari, and F. Babiloni, "Multiple pathways analysis of brain functional networks from EEG signals: an application to real data," *Brain Topogr*, vol. 23, no. 4, pp. 344–354, January 2011.
- [161] A. Zalesky, L. Cocchi, A. Fornito, M. M. Murray, and E. Bullmore, "Connectivity differences in brain networks," *Neuroimage*, vol. 60, no. 2, pp. 1055–1062, April 2012.
- [162] C. E. Schroeder and P. Lakatos, "The gamma oscillation: master or slave?" *Brain Topogr*, vol. 22, no. 1, pp. 24–26, June 2009.
- [163] L. Uldry, "Therapeutic strategies for the treatment of atrial fibrillation: new insights from biophysical modeling and signal processing," Ph.D. dissertation, École polytechnique fédérale de Lausanne, 2011, director: J.-M. Vesin.
- [164] E. Braunwald, Ed., *Heart disease: a textbook of cardiovascular medicine*, 4th ed. W. B. Saunders, 1992.
- [165] A. J. Camm, P. Kirchhof, G. Y. Lip, U. Schotten, I. Savelieva, S. Ernst, I. C. Van Gelder, N. Al-Attar, G. Hindricks, B. Prendergast, H. Heidbuchel, O. Alfieri, A. Angelini, D. Atar, P. Colonna, R. De Caterina, J. De Sutter, A. Goette, B. Gorenek, M. Heldal, S. H. Hohloser, P. Kolh, J.-Y. Le Heuzey, P. Ponikowski, and F. H. Rutten, "Guidelines for the management of atrial fibrillation," *Eur Heart J*, vol. 31, no. 19, pp. 2369–2429, October 2010.

- [166] L. S. Wann, A. B. Curtis, C. T. January, K. A. Ellenbogen, J. E. Lowe, N. M. Estes, R. L. Page, M. D. Ezekowitz, D. J. Slotwiner, W. M. Jackman, W. G. Stevenson, and C. M. Tracy, "2011 ACCF/AHA/HRS focused update on the management of patients with atrial fibrillation (updating the 2006 guideline)," *Circulation*, vol. 123, no. 1, pp. 104–123, January 2011.
- [167] V. Fuster, L. E. Rydén, D. S. Cannom, H. J. Crijns, A. B. Curtis, K. A. Ellenbogen, J. L. Halperin, G. N. Kay, J.-Y. Le Huezey, J. E. Lowe, S. B. Olsson, E. N. Prystowsky, J. L. Tamargo, and L. S. Wann, "2011 ACCF/AHA/HRS focused updates incorporated into the ACC/AHA/ESC 2006 guidelines for the management of patients with atrial fibrillation: a report of the American College of Cardiology Foundation/American Heart Association Task Force on practice guidelines developed in partnership with the European Society of Cardiology and in collaboration with the European Heart Rhythm Association and the Heart Rhythm Society," *J Am Coll Cardiol*, vol. 57, no. 11, pp. e101–e198, March 2011.
- [168] F. Jung and J. P. DiMarco, "Treatment strategies for atrial fibrillation," *Am J Med*, vol. 104, no. 3, pp. 272–286, March 1998.
- [169] G. Thrall, D. Lane, D. Carroll, and G. Y. H. Lip, "Quality of life in patients with atrial fibrillation: a systematic review," *Am J Med*, vol. 119, no. 5, pp. 448.e1–448.e19, May 2006.
- [170] E. J. Benjamin, P. A. Wolf, R. B. D'Agostino, H. Silbershatz, W. B. Kannel, and D. Levy, "Impact of atrial fibrillation on the risk of death: the Framingham heart study," *Circulation*, vol. 98, no. 10, pp. 946–952, September 1998.
- [171] Y. Miyasaka, M. E. Barnes, B. J. Gersh, S. S. Cha, K. R. Bailey, W. P. Abhayaratna, J. B. Seward, and T. S. Tsang, "Secular trends in incidence of atrial fibrillation in Olmsted County, Minnesota, 1980 to 2000, and implications on the projections for future prevalence," *Circulation*, vol. 114, no. 2, pp. 119–125, July 2006.
- [172] J. S. Steinberg, "Atrial fibrillation: an emerging epidemic?" *Heart*, vol. 90, no. 3, pp. 239–240, March 2004.
- [173] G. Y. H. Lip and H.-F. Tse, "Management of atrial fibrillation," *Lancet*, vol. 370, no. 9587, pp. 604–618, August 2007.
- [174] M. C. E. F. Wijffels, C. J. H. J. Kirchhof, R. Dorland, and M. A. Allessie, "Atrial fibrillation begets atrial fibrillation: a study in awake chronically instrumented goats," *Circulation*, vol. 92, no. 7, pp. 1954–1968, October 1995.
- [175] M. Lemay, J.-M. Vesin, A. van Oosterom, V. Jacquemet, and L. Kappenberger, "Cancellation of ventricular activity in the ECG: evaluation of novel and existing techniques," *IEEE T Bio-Med Eng*, vol. 54, no. 3, pp. 542–546, March 2007.
- [176] A. Bollmann, D. Husser, L. Mainardi, F. Lombardi, P. Langley, A. Murray, J. J. Rieta, J. Millet, S. B. Olsson, M. Stridh, and L. Sörnmo, "Analysis of surface electrocardiograms in atrial fibrillation: techniques, research, and clinical applications," *Europace*, vol. 8, no. 11, pp. 911–926, November 2006.
- [177] S. Matsuo, N. Lellouche, M. Wright, M. Bevilacqua, S. Knecht, I. Nault, K.-T. Lim, L. Arantes, M. D. O'Neill, P. G. Platonov, J. Carlson, F. Sacher, M. Hocini, P. Jaïs, and M. Haïssaguerre, "Clinical predictors of termination and clinical outcome of catheter ablation for persistent atrial fibrillation," *J Am Coll Cardiol*, vol. 54, no. 9, pp. 788–795, August 2009.

- [178] T. Rostock, T. V. Salukhe, D. Steven, I. Drewitz, B. A. Hoffmann, K. Bock, H. Servatius, K. Müllerleile, A. Sultan, N. Gosau, T. Meinertz, K. Wegscheider, and S. Willems, "Long-term single- and multiple-procedure outcome and predictors of success after catheter ablation for persistent atrial fibrillation," *Heart Rhythm*, vol. 8, no. 9, pp. 1391–1397, September 2011.
- [179] I. Nault, N. Lellouche, S. Matsuo, S. Knecht, M. Wright, K.-T. Lim, F. Sacher, P. Platonov, A. Deplagne, P. Bordachar, N. Derval, M. D. O'Neill, G. J. Klein, M. Hocini, P. Jaïs, J. Clémenty, and M. Haïssaguerre, "Clinical value of fibrillatory wave amplitude on surface ECG in patients with persistent atrial fibrillation," *J Interv Card Electr*, vol. 26, no. 1, pp. 11–19, October 2009.
- [180] M. Meo, V. Zarzoso, O. Meste, D. G. Latcu, and N. Saoudi, "Non-invasive prediction of catheter ablation outcome in persistent atrial fibrillation by exploiting the spatial diversity of surface ECG," in *Conf Proc IEEE Eng Med Biol Soc*, August 2011, pp. 5531–5534.
- [181] T. H. Everett, J. R. Moorman, L.-C. Kok, J. G. Akar, and D. E. Haines, "Assessment of global atrial fibrillation organization to optimize timing of atrial defibrillation," *Circulation*, vol. 103, no. 23, pp. 2857–2861, June 2001.
- [182] L. Uldry, J. Van Zaen, L. Kappenberger, and J.-M. Vesin, "Novel complexity measurements in the ECG differentiate persistent from permanent atrial fibrillation," *Europace*, vol. 13, no. suppl 3, p. P511, June 2011.
- [183] L. Uldry, J. Van Zaen, Y. Prudat, L. Kappenberger, and J.-M. Vesin, "Measures of spatiotemporal organization differentiate persistent from long-standing atrial fibrillation," *Europace*, vol. 14, no. 8, pp. 1125–1131, August 2012.
- [184] M. Stridh, D. Husser, A. Bollmann, and L. Sörnmo, "Waveform characterization of atrial fibrillation using phase information," *IEEE T Bio-Med Eng*, vol. 56, no. 4, pp. 1081–1089, April 2009.
- [185] Y. Takahashi, P. Sanders, P. Jaïs, M. Hocini, R. Dubois, M. Rotter, T. Rostock, C. J. Nalliah, F. Sacher, J. Clémenty, and M. Haïssaguerre, "Organization of frequency spectra of atrial fibrillation: relevance to radiofrequency catheter ablation," *J Cardiovasc Electr*, vol. 17, no. 4, pp. 382–388, April 2006.
- [186] M. Holm, S. Pehrson, M. Ingemansson, L. Sörnmo, R. Johansson, L. Sandhall, M. Sune-mark, B. Smideberg, C. Olsson, and S. B. Olsson, "Non-invasive assessment of the atrial cycle length during atrial fibrillation in man: introducing, validating and illustrating a new ECG method," *Cardiovasc Res*, vol. 38, no. 1, pp. 69–81, April 1998.
- [187] K. Knight, *Mathematical statistics*. Chapman & Hall, 2000.
- [188] Y. Hochberg and A. C. Tamhane, *Multiple comparison procedures*. John Wiley & Sons, 1987.
- [189] P. G. Platonov, I. Nault, F. Holmqvist, M. Stridh, M. Hocini, and M. Haïssaguerre, "Left atrial appendage activity translation in the standard 12-lead ECG," *J Cardiovasc Electr*, vol. 22, no. 6, pp. 706–710, June 2011.
- [190] S. Petrutiu, A. V. Sahakian, W. Fisher, and S. Swiryn, "Manifestation of left atrial events and interatrial frequency gradients in the surface electrocardiogram during atrial fibrillation: contributions from posterior leads," *J Cardiovasc Electr*, vol. 20, no. 11, pp. 1231–1236, November 2009.

- [191] M. Haïssaguerre, P. Sanders, M. Hocini, Y. Takahashi, M. Rotter, F. Sacher, T. Rostock, L.-F. Hsu, P. Bordachar, S. Reuter, R. Roudaut, J. Clémenty, and P. Jaïs, "Catheter ablation of long-lasting persistent atrial fibrillation: critical structures for termination," *J Cardiovasc Electr*, vol. 16, no. 11, pp. 1125–1137, November 2005.
- [192] B. Widrow and S. D. Stearns, *Adaptive signal processing*. Prentice-Hall, 1985.
- [193] S. Racinais, D. Bishop, R. Denis, G. Lattier, A. Mendez-Villaneuva, and S. Perrey, "Muscle deoxygenation and neural drive to the muscle during repeated sprint cycling," *Med Sci Sport Exer*, vol. 39, no. 2, pp. 268–274, February 2007.
- [194] M. Buchheit, C. R. Abbiss, J. J. Peiffer, and P. B. Laursen, "Performance and physiological responses during a sprint interval training session: relationships with muscle oxygenation and pulmonary oxygen uptake kinetics," *Eur J Appl Physiol*, vol. 112, no. 2, pp. 767–779, February 2012.
- [195] R. Faiss, B. Léger, P. Fournier, Y. Eggel, O. Dériaz, and G. P. Millet, "Repeated-sprint ability is further enhanced by intensive training in hypoxia than in normoxia," in *Proc Physiol Soc*, vol. 26, March 2012, p. PC4.
- [196] F. Cottin, C. Médigue, P. Lopes, P.-M. Leprêtre, R. Heubert, and V. Billat, "Ventilatory thresholds assessment from heart rate variability during an incremental exhaustive running test," *Int J Sports Med*, vol. 28, no. 4, pp. 287–294, April 2007.
- [197] G. K. Karapetian, H. J. Engels, and R. J. Gretebeck, "Use of heart rate variability to estimate LT and VT," *Int J Sports Med*, vol. 29, no. 8, pp. 652–657, August 2008.
- [198] O. Meste, G. Blain, and B. S., "Analysis of the respiratory and cardiac systems coupling in pyramidal exercise using a time-varying model," in *Comput Cardiol*, vol. 29, September 2002, pp. 429–432.
- [199] G. Blain, O. Meste, T. Bouchard, and S. Bermon, "Assessment of ventilatory thresholds during graded and maximal exercise test using time varying analysis of respiratory sinus arrhythmia," *Brit J Sport Med*, vol. 39, no. 7, pp. 448–452, July 2005.
- [200] T. K. Moon and W. C. Stirling, *Mathematical methods and algorithms for signal processing*. Prentice Hall, 2000.
- [201] I. S. Gradshteyn and I. M. Ryzhik, *Table of integrals, series, and products*, 7th ed., A. Jeffrey and D. Zwillinger, Eds. Academic Press, 2007.
- [202] M. Abramowitz and I. A. Stegun, *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*, 8th ed. Dover Publications, 1973.
- [203] A. Papoulis and S. U. Pillai, *Probability, random variables and stochastic processes*, 4th ed. McGraw-Hill, 2002.
- [204] J. V. Michalowicz, J. M. Nichols, F. Bucholtz, and C. C. Olson, "An Isserlis' theorem for mixed Gaussian variables: application to the auto-bispectral density," *J Stat Phys*, vol. 136, no. 1, pp. 89–102, July 2009.
- [205] G. C. Wick, "The evaluation of the collision matrix," *Phys Rev*, vol. 80, no. 2, pp. 268–272, October 1950.
- [206] L. Isserlis, "On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables," *Biometrika*, vol. 12, no. 1/2, pp. 134–139, November 1918.

- [207] I. S. Reed, "On a moment theorem for complex Gaussian processes," *IRE T Inform Theor*, vol. 8, no. 3, pp. 194–195, April 1962.
- [208] A. Jonquière, "Note sur la série $\sum_{n=1}^{n=\infty} \frac{x^n}{n^s}$," *B Soc Math Fr*, vol. 17, pp. 142–152, April 1889.
- [209] L. Lewin, *Polylogarithms and associated functions*. North Holland, 1981.

Index

- adaptive line enhancer, 16–18
- amplitude modulation, 9
- analytic signal, 9–10
- atrial fibrillation, 117

- catheter ablation, 118
- complex adaptive notch filter, 21
- continuous energy separation algorithm, 15
- cross-frequency couplings, 92

- discrete energy separation algorithm, 15
- discrete Fourier transform, *see* Fourier transform
- discrete oscillator equation
 - complex discrete oscillator equation, 32, 76
 - real discrete oscillator equation, 18
- discrete-time Fourier transform, *see* Fourier transform

- electrocardiogram, 117
- electroencephalogram, 91

- Fourier transform, 145
 - discrete Fourier transform, 145
 - discrete-time Fourier transform, 145
- frequency modulation, 9

- generalized adaptive notch filter, 19–20

- harmonic frequency tracker, *see* single frequency tracker
- Hilbert transform, 10
 - discrete Hilbert transform, 11–13

- instantaneous frequency, 9–11

- multiple frequency tracker, *see* single frequency tracker
- multivariate frequency tracker, *see* single frequency tracker
- multivariate multiple frequency tracker, *see* single frequency tracker

- near-infrared spectroscopy, 132
- non-causal frequency estimator, 75–78

- oscillator-based adaptive band-pass filter, 18–19

- photoplethysmogram, 131
- power spectral density, 1

- R-R interval, 85, 135
- robust modified Newton algorithm, 22–24

- single frequency tracker, 31–33, 147
 - harmonic frequency tracker, 69–70
 - multiple frequency tracker, 58–59
 - multivariate frequency tracker, 60–63
 - multivariate multiple frequency tracker, 66–67
- surrogate data, 95

Jérôme VAN ZAEN

Rue Louis Meyer 5
CH-1800 Vevey
Phone: +41 78 772 96 47
Email: jerome.vanzaen@gmail.com

Date of birth: August 14, 1983
Nationality: Swiss/Belgian
Marital status: Single

Education

- 2008 – 2012** ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE, Switzerland.
Ph.D. in Biomedical Signal Processing under the direction of Dr. J.-M. Vesin.
Thesis entitled *Efficient schemes for adaptive frequency tracking and their relevance for EEG and ECG*.
- 2002 – 2007** ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE, Switzerland.
M.Sc. in Communication Systems.
Specialization : Signals and images.
- 1999 – 2002** GYMNASSE DE CHAMBLANDES, Pully, Switzerland.
Federal Diploma of Maturity — scientific section.

Research Experience

- 2012 – 2013** CENTRE HOSPITALIER UNIVERSITAIRE VAUDOIS, Lausanne, Switzerland.
Civilian service as a research assistant in Radiology Service.
- 2008 – 2012** ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE, Switzerland.
Research assistant in Applied Signal Processing Group.
- 2007** ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE, Switzerland.
Master thesis entitled *Development of adaptive frequency tracking algorithms* at Signal Processing Institute.

Teaching Experience

- 2010 – 2011** ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE, Switzerland.
Teaching assistant for lab sessions in the course *Biomedical signal processing*.
- 2009** ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE, Switzerland.
Teaching assistant for lab sessions in the course *Analyse informatique de données*.

Language

- French : native language.
English : advanced.

Computer Skills

- Environment : Linux/Unix and Windows.
Programming : Matlab, wide experience.
Python, good knowledge.
C/C++ and Java, basic knowledge.
Tools : Office, OpenOffice and LaTeX.

Interests and Activities

- Reading.
- Running, hiking.

References

Available upon request.

Publications

Journal Papers

F. Degache, J. Van Zaen, L. Oehen, K. Guex, F. Schena, and G. P. Millet, “Effects of sleep deprivation combined with muscle fatigue on somatosensory integration: posturographic study,” *Chronobiology International*, 2012, in preparation.

J. Van Zaen, M. M. Murray, R. A. Meuli, and J.-M. Vesin, “Adaptive filtering methods for identifying cross-frequency couplings in human EEG,” *PLoS One*, 2012, submitted.

A. Buttu, J. Van Zaen, A. Viso, A. Forclaz, P. Pascale, S. M. Narayan, E. Pruvot, and J.-M. Vesin, “Adaptive frequency tracking of the baseline ECG identifies the site of atrial fibrillation termination by catheter ablation,” *Biomedical Signal Processing and Control*, 2012, submitted.

L. Uldry, J. Van Zaen, Y. Prudat, L. Kappenberger, and J.-M. Vesin, “Measures of spatiotemporal organization differentiate persistent from long-standing atrial fibrillation,” *Europace*, vol. 14, no 8, pp. 1125–1131, August 2012.

J. Van Zaen, L. Uldry, C. Duchêne, Y. Prudat, R. A. Meuli, M. M. Murray, and J.-M. Vesin, “Adaptive tracking of EEG oscillations,” *Journal of Neuroscience Methods*, vol. 186, no. 1, pp. 97–106, January 2010.

Conference Papers and Abstracts

A. Buttu, J. Van Zaen, A. Viso, A. Forclaz, P. Pascale, S. M. Narayan, J.-M. Vesin, and E. Pruvot, “Termination of atrial fibrillation by catheter ablation can be successfully predicted from baseline ECG.,” *Computing in Cardiology*, Kraków, Poland, September 2012.

A. Buttu, J. Van Zaen, A. Viso, A. Forclaz, P. Pascale, S. M. Narayan, J.-M. Vesin, and E. Pruvot, “Termination of atrial fibrillation by catheter ablation can be successfully predicted from baseline ECG,” *Cardiostim*, Nice, France, June 2012.

A. Buttu, J. Van Zaen, A. Viso, A. Forclaz, P. Pascale, S. M. Narayan, J.-M. Vesin, and E. Pruvot, “Termination of atrial fibrillation by catheter ablation can be successfully predicted from baseline ECG,” *Swiss Society of Cardiology Annual Meeting*, Lausanne, Switzerland, June 2012.

A. Buttu, J. Van Zaen, A. Viso, A. Forclaz, P. Pascale, S. M. Narayan, J.-M. Vesin, and E. Pruvot, "A new method for ECG tracking of persistent atrial fibrillation termination during stepwise ablation," *Computing in Cardiology*, Hangzhou, China, September 2011.

J. Van Zaen, L. Uldry, A. Buttu, L. Kappenberger, and J.-M. Vesin, "Harmonic frequency tracking algorithm for predicting the success of pharmacological cardioversion of atrial fibrillation," *EHRA Europace*, Madrid, Spain, June 2011.

L. Uldry, J. Van Zaen, L. Kappenberger, and J.-M. Vesin, "Novel complexity measurements in the ECG differentiate persistent from permanent atrial fibrillation," *EHRA Europace*, Madrid, Spain, June 2011.

J. Van Zaen, M. M. Murray, R. A. Meuli, and J.-M. Vesin, "Cross-frequency coupling in EEG revealed by adaptive oscillation tracking," *16th Annual Meeting of the Organization for Human Brain Mapping*, Barcelona, Spain, June 2010.

M. Jabloun, J. Van Zaen, and J.-M. Vesin, "Time-frequency analysis based on the phase-rectified signal averaging method," *17th European Signal Processing Conference*, Glasgow, Scotland, August 2009.

J. Van Zaen, L. Uldry, J.-M. Vesin, M. M. Murray, and R. A. Meuli, "Tracking of oscillatory components in EEG," *Swiss Society for Neuroscience Annual Meeting*, Fribourg, Switzerland, March 2009.

Y. Prudat, J. Van Zaen, and J.-M. Vesin, "Harmonic extension of an adaptive notch filter for frequency tracking," *16th European Signal Processing Conference*, Lausanne, Switzerland, August 2008.

