

Kalman filter methods for real-time signal processing in tokamak fusion research

Diogo Miguel Louro Alves

Supervisor: Doctor Rui Miguel Dias Alves Coelho

Co-Supervisor: Doctor Bernardo Brotas de Carvalho

Thesis approved in public session to obtain the PhD degree in

Technological Physics Engineering

Jury final classification: Pass with merit

Jury

Chairperson: Chairman of the IST Scientific Board

Members of the committee:

Doctor Horácio João Matos Fernandes
Doctor João Miguel Raposo Sanches
Doctor Francisco José Baptista Salzedas
Doctor Rui Miguel Dias Alves Coelho
Doctor Andrea Murari

Kalman filter methods for real-time signal processing in tokamak fusion research

Diogo Miguel Louro Alves

Supervisor: Doctor Rui Miguel Dias Alves Coelho

Co-Supervisor: Doctor Bernardo Brotas de Carvalho

Thesis approved in public session to obtain the PhD degree in

Technological Physics Engineering

Jury final classification: Pass with merit

Jury

Chairperson: Chairman of the IST Scientific Board

Members of the committee:

Doctor Horácio João Matos Fernandes, Professor Auxiliar (com Agregação) do Instituto Superior Técnico, da Universidade de Lisboa;

Doctor João Miguel Raposo Sanches, Professor Auxiliar (com Agregação) do Instituto Superior Técnico, da Universidade de Lisboa;

Doctor Francisco José Baptista Salzedas, Professor Auxiliar da Faculdade de Engenharia da Universidade do Porto;

Doctor Rui Miguel Dias Alves Coelho, Investigador Auxiliar Convidado do Instituto Superior Técnico, da Universidade de Lisboa;

Doctor Andrea Murari, Investigador do Consorzio RFX Padova, Italia.

Resumo

Esta dissertação descreve vários métodos de processamento de sinal baseados em filtros de Kalman e desenvolvidos para resolver problemas específicos na investigação em máquinas de fusão do tipo tokamak. A relevância de métodos ótimos de estimativa de estado, e do filtro de Kalman em particular, cuja competência é amplamente reconhecida, é apresentada como uma alternativa viável a técnicas mais tradicionais de processamento de sinal. Após a descrição do enquadramento e motivação para o uso de técnicas de filtragem, o filtro de Kalman é derivado como a solução ótima do problema de estimação linear do estado de um sistema no sentido da minimização da média dos quadrados do erro. As mesmas equações previamente obtidas são de novo derivadas, desta feita assumindo distribuições Gaussianas de ruído no processo e na medida, satisfazendo o critério do *máximo a posteriori* no contexto da análise Bayesiana.

Vários métodos baseados em implementações lineares do filtro de Kalman são propostos e alguns resultados são apresentados no contexto da: (a) estimativa da amplitude de componentes de frequência quase-periódicas para processamento de sinal em tempo-real de dados do diagnóstico de efeito de Stark cinético no tokamak Joint European Torus (JET); (b) a análise do espectro de coerência para determinação em tempo-real da periodicidade espacial de actividade MagnetoHidrodinâmica (MHD) no tokamak ISTTOK; (c) a detecção síncrona de sinal necessária no processamento de dados do diagnóstico da antena dos modos próprios toroidais de Alfvén no tokamak JET; (d) a estimativa da corrente no circuito RL usado para o controlo das perturbações magnéticas externas no tokamak JET.

Uma outra implementação não-linear do filtro de Kalman para rastreamento de componentes em frequência é explorada e estudada no contexto da estimativa da amplitude, frequência e fase instantâneas de componentes não estacionárias. É também apresentado um esquema adaptativo para minimizar os efeitos da contaminação inter-frequências e perda de rastreamento encontrados não só em estimativas de uma só componente mas também na estimativa simultânea de várias componentes. Um algoritmo híbrido baseado neste mecanismo e também numa implementação linear do filtro de Kalman foi igualmente desenvolvido com sucesso para a análise em tempo-real da amplitude, frequência e número de modo da actividade MHD no tokamak JET. Finalmente, é também descrita a implementação de uma aplicação em tempo-real, que inclui um filtro de Kalman linear, para o controlo de perturbações magnéticas externas no tokamak JET.

Abstract

This dissertation describes a number of Kalman filter-based signal processing methods developed to address dedicated problems encountered in tokamak fusion research. The relevance of optimal state estimation methods and the Kalman filter in particular, whose competence is widely recognized, is put forward as a viable alternative to other, more traditional, signal processing approaches. Some general background and motivation for the use of dedicated filtering techniques is introduced and the Kalman filter is derived as the optimal solution of the linear state estimation problem in the minimum mean square error sense. Furthermore, the same Kalman filter equations are derived under Gaussian assumptions and satisfying the maximum a posteriori criterion in the context of Bayesian analysis.

Methods based on linear Kalman filter implementations are proposed and results shown in the context of: (a) the amplitude estimation of quasi-periodic frequency components in the real-time signal processing of data from the motional Stark effect diagnostic at the Joint European Torus (JET) tokamak; (b) the coherence spectrum analysis for real-time mode number estimation of MagnetoHyDrodynamics (MHD) activity in the ISTTOK tokamak; (c) the synchronous detection required by the toroidal Alfvén eigenmode (TAE) antenna diagnostic at JET; (d) the estimation of the current flowing in the RL circuit used for the feedback control of external magnetic perturbations in the JET tokamak.

Also, a non-linear frequency tracking implementation of the Kalman filter is explored and thoroughly studied for the real-time estimation of the instantaneous amplitude, frequency and phase of non-stationary signal components. An adaptive scheme is developed for coping with inter-tone contamination and tracking loss issues encountered and its competences are demonstrated not only in single-tone but also in simultaneous multi-tone estimations. A hybrid algorithm based on the previous together with a linear Kalman filter implementation has also been successfully developed for the real-time analysis of the amplitude, frequency and mode number estimation of (MHD) activity in the JET tokamak. Finally, the feedback control application, which includes a linear Kalman filter implementation, used for the control of external magnetic perturbations in the JET tokamak is thoroughly described.

Keywords

Nuclear fusion
Tokamak
Joint European Torus
Digital signal processing
Real-time systems
Control
Kalman filter
State estimation
Harmonic estimation
Frequency tracking
Magnetohydrodynamics

Palavras-chave

Fusão nuclear
Tokamak
Joint European Torus
Processamento digital de sinal
Sistemas de tempo-real
Controlo
Filtro de Kalman
Estimativa de estado
Estimativa de harmónicas
Rastreamento de frequências
Magnetohidrodinâmica

Acknowledgements

I would like to thank Prof. F. Serra, Prof. C. Varandas, Prof. H. Fernandes, Dr. B. Gonçalves, Dr. P. Lomas, Dr. F. Sartori, Dr. F. Rimini and Dr. A. Neto for making possible the professional path that lead to the present dissertation.

I would also like to express my gratitude to my colleague Dr. D. Valcarcel and especially to Dr. A. Neto, a key person in carving my professional path, for his never ending support, professional example, the numerous and stimulating discussions and above all for his friendship.

I acknowledge the availability and kind assistance of Mrs. C. Nunes and Eng. J. Santos in the assembly of the documents required for the submission of this dissertation.

Finally I'd like to thank Dr. R. Coelho, my supervisor who, over the years, with infinite patience, kindness and availability, taught me and steered my scientific path towards this dissertation.

Last and most importantly I would like to thank my family for the unconditional love, and in many occasions sacrifice, making this thesis possible.

This work was supported by EURATOM and carried out within the framework of the European Fusion Development Agreement. IST activities also received financial support from “Fundação para a Ciência e Tecnologia” through project Pest-OE/SADG/LA0010/2011. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

*To Elsa, Beatriz, Maria and
my parents. With love.*

Contents

1	Introduction	1
1.1	Nuclear fusion power	1
1.2	Tokamaks	2
1.3	Real-time systems in the world's largest tokamak	3
1.4	Real-time signal processing	6
2	Kalman-based state estimation	9
2.1	Brief introduction to the filtering problem	9
2.2	Kalman filter	10
2.2.1	Minimum mean square error derivation	10
2.2.2	Bayesian derivation	15
2.2.3	Scalar time-invariant Kalman filter gain	18
2.3	Non-linear Kalman filters	20
2.3.1	Extended Kalman filter	20
2.3.2	Iterated extended Kalman filter	21
2.3.3	Robust extended Kalman filter	22
2.3.4	Unscented Kalman filter	22
3	Linear harmonic estimation	25
3.1	Introduction	25
3.2	Theory	26
3.3	Applications	27
3.3.1	MSE signal processing	27
3.3.2	Coherence spectrum	42
3.3.3	TAE signal processing I	47
4	Non-stationary and non-linear frequency tracking and harmonic estimation	55
4.1	Introduction	55
4.2	Theory	55
4.3	Studies	57
4.3.1	Single Tone	57
4.3.2	Amplitude Separation	58
4.3.3	Frequency Separation	61
4.3.4	Multi-Tone Harmonic Estimation and Frequency Tracking	63
4.3.5	Remarks	63
4.4	Applications	64
4.4.1	TAE signal processing II	64

4.4.2	A real-time algorithm for the harmonic estimation and frequency tracking of dominant signal components in fusion plasma magnetic diagnostics	65
4.4.3	Kalman filter methods for real-time frequency and mode number estimation of MHD activity in tokamak plasmas	79
5	Control of external magnetic perturbations in the JET tokamak	99
5.1	Introduction	99
5.1.1	Overview of the EFCC current controller system	100
5.1.2	User configuration	101
5.2	Hardware	102
5.2.1	Error Field Correction Coils	102
5.2.2	Voltage Controlled Voltage Sources	102
5.2.3	Controller system	102
5.3	Software	103
5.3.1	Multithreaded Application Real-Time executor	103
5.3.2	Synchronization	104
5.3.3	Control modules	105
5.3.4	Control modes	112
5.4	Preliminary results	118
5.4.1	Current reference multiplexing logic assessment	119
5.4.2	VCVS model performance	120
5.4.3	Controller performace	121
6	Final remarks and future work	125
6.1	Future work	129
	Bibliography	131

List of Figures

1.1	Effective collision cross-section of different types of fusion reactions. <i>Courtesy of EFDA-JET</i>	2
1.2	Particle confinement. <i>Courtesy of EFDA-JET</i>	3
1.3	Tokamak sketch. <i>Courtesy of EFDA-JET</i>	4
1.4	Tokamak transformer. <i>Courtesy of EFDA-JET</i>	4
2.1	Steady-state time-invariant scalar Kalman filter gain contour as a function of R and Q	19
3.1	Overview of the MSE diagnostic apparatus at JET. <i>Courtesy of EFDA-JET</i>	28
3.2	Simplified single channel light path for the MSE diagnostic at JET. <i>Courtesy of EFDA-JET</i>	29
3.3	JET pulse #85394 - (a) Plasma current; (b) Toroidal field; (c) NBI power; (d) Vertical field B_z and radial position of the magnetic axis (black curve) as calculated by an off-line equilibrium reconstruction code.	31
3.4	Frequency response of a 15 sample square window using the DFT alone (black line) and jointly with the KF, where the r/q ratios are equal to 10^{-4} (blue), 10^1 (red), and 10^2 (green).	33
3.5	(a) Amplitude and (b) phase response of the low-pass estimator; (c) Amplitude and (d) phase response of the band-pass estimator.	35
3.6	Phase response zoom of the band-pass estimator.	36
3.7	Effect of a frequency mismatch on the estimated amplitude of an AM signal ($q = 10^{-4}$). (a) Offset in the mean amplitude estimation increases with covariance r and mismatch frequency $\Delta\omega/2\pi$. (b) (Orange) Time evolution of the estimated amplitude for the $r = 1$ and 100 Hz mismatch case is shown.	37
3.8	Spectrum of channel 15's APD signal for JET pulse #70129 indicating which of the frequency tones are used in KFHE's state transition model.	39
3.9	(a) D_α radiated power at the outer divertor leg; (b) APD signal of channel 15 (black curve) and DC component estimation of the multi-tone KFHE (gray curve).	40
3.10	Amplitude of the 46 kHz component using a lock-in amplifier (black) and the Kalman filter algorithm (red) with ELM+50 Hz removal. The ratio r/q , set to 10^5 in between ELMs, is reduced by three orders of magnitude during the ELMs	41
3.11	Spectrum of the 46 kHz component using a the SPL (black) and the KF-based method (red) with the ELM+50 Hz removal controller disabled	42
3.12	Radial profile of the pitch angle calculation using the SPL (black) and the KF-based method (gray) at: (a) $t \sim 61.0825$ s; (b) $t \sim 61.0845$ s and (c) $t \sim 61.0925$ s	43

3.13	Mirnov coil signal ($\theta = 60^\circ$, pulse #17081) evidencing an $m = 2$ mode rotating at ~ 160 kHz	46
3.14	KFHE-based - (a) Coherency spectrum and (b) Poloidal mode number spectrum for $\gamma > 0.8$	46
3.15	Fourier-based - (a) Coherency spectrum and (b) Poloidal mode number spectrum for $\gamma > 0.95$	47
3.16	Typical spectrogram of a Mirnov coil signal showing clearly the frequency sweeps and resonant spots	48
3.17	The DPL estimate (blue), KFHE estimate (red) and true value (black) of: (a) amplitude of the plasma response and (b) phase difference between excitation and plasma response.	51
3.18	Frequency evolution of the excitation waveform.	52
3.19	Amplitude estimation - DPL estimate with 20 samples for LPF-ing (blue), KFHE estimate (red) and true amplitude (black).	53
4.1	Estimation of a non-stationary single-tone in a noisy signal	58
4.2	Estimation of the frequency of the 15 kHz component for the cases where a_2 is 10^0 (blue), 10^1 (green), 10^2 (red) and 5×10^2 (cyan).	59
4.3	Estimates of the 15 kHz signal component in the cases where a_2 is 10^0 (blue), 10^1 (green), 10^2 (red) and 5×10^2 (cyan).	60
4.4	Estimates of the 20 kHz signal component in the cases where a_2 is 10^0 (blue), 10^1 (green), 10^2 (red) and 5×10^2 (cyan).	61
4.5	Estimates of two signal components of equal amplitudes.	62
4.6	Estimates of two signal components of different amplitudes in a noisy signal.	62
4.7	Multi-tone estimation in a noisy signal.	63
4.8	Amplitude estimation performance comparison between the EKF (blue) and KF (red).	64
4.9	EKF frequency estimation (blue) and true frequency content (red).	65
4.10	Top plot: raw signal spectrogram; 2nd plot: estimated frequency; Bottom plot: estimated amplitude. Blue: EKFFT; Green: rEKFFT; Red: UKFFT; Cyan: DFT.	67
4.11	JET pulse 68628 - Top plot: raw signal spectrogram (\log_{10} scale); 2nd plot: estimated frequency; 3rd plot: normalized frequency estimation error; 4th plot: estimated amplitude; 5th plot: normalized amplitude estimation error; Bottom plot: 10 ms averaged residue. Blue: EKFFT; Green: rEKFFT; Red: UKFFT; Black: DFT.	69
4.12	JET pulse 68628 - Normalized amplitude inter-frequency contamination of the estimated signal - Top plot: EKFFT; Middle plot: rEKFFT; Bottom plot: UKFFT.	70
4.13	JET pulse 68628 - Adaptive estimation - Top plot: raw signal spectrogram (\log_{10} scale); 2nd plot: estimated frequency; 3rd plot: normalized frequency estimation error; 4th plot: estimated amplitude; 5th plot: normalized amplitude estimation error; Bottom plot: 10 ms averaged residue. Blue: EKFFT; Green: rEKFFT; Red: UKFFT; Black: DFT.	71
4.14	JET pulse 68628 - Adaptive estimation - Normalized amplitude inter-frequency contamination of the estimated signal - Top plot: EKFFT; Middle plot: rEKFFT; Bottom plot: UKFFT.	72

4.15	JET pulse 67694 - Adaptive estimation - Top plot : raw signal spectrogram (\log_{10} scale); 2nd plot : estimated frequency; 3rd plot : normalized frequency estimation error; 4th plot : estimated amplitude; 5th plot : normalized amplitude estimation error; Bottom plot : 10 ms averaged residue. Blue : EKFFT; Green : rEKFFT; Red : UKFFT; Black : DFT.	72
4.16	JET pulse 67694 - Adaptive estimation - Normalized amplitude inter-frequency contamination of the estimated signal - Top plot : EKFFT; Middle plot : rEKFFT; Bottom plot : UKFFT.	73
4.17	JET pulse 79743 - Adaptive estimation - Top plot : raw signal spectrogram (\log_{10} scale); 2nd plot : estimated frequency; 3rd plot : normalized frequency estimation error; 4th plot : estimated amplitude; 5th plot : normalized amplitude estimation error; Bottom plot : 10 ms averaged residue. Blue : EKFFT; Green : rEKFFT; Red : UKFFT; Black : DFT.	73
4.18	JET pulse 79743 - Adaptive estimation - Normalized amplitude inter-frequency contamination of the estimated signal - Top plot : EKFFT; Middle plot : rEKFFT; Bottom plot : UKFFT.	74
4.19	JET pulse 68628 - Three tone adaptive EKFFT - Top plot : raw signal spectrogram (\log_{10} scale); 2nd plot : estimated frequencies; 3rd plot : amplitude of tone a); 4th plot : amplitude of tone b); 5th plot : amplitude of tone c); Bottom plot : residues. Red : tone a) estimate; Green : tone b) estimate; Blue : tone c) estimate; Black : DFT.	76
4.20	JET pulse 68628 - Three tone adaptive rEKFFT - Top plot : raw signal spectrogram (\log_{10} scale); 2nd plot : estimated frequencies; 3rd plot : amplitude of tone a); 4th plot : amplitude of tone b); 5th plot : amplitude of tone c); Bottom plot : residues. Red : tone a) estimate; Green : tone b) estimate; Blue : tone c) estimate; Black : DFT.	77
4.21	JET pulse 68628 - Three tone adaptive UKFFT - Top plot : raw signal spectrogram (\log_{10} scale); 2nd plot : estimated frequencies; 3rd plot : amplitude of tone a); 4th plot : amplitude of tone b); 5th plot : amplitude of tone c); Bottom plot : residues. Red : tone a) estimate; Green : tone b) estimate; Blue : tone c) estimate; Black : DFT.	77
4.22	JET pulse 68558 - Three tone adaptive EKFFT - Top plot : raw signal spectrogram (\log_{10} scale); 2nd plot : estimated frequencies; 3rd plot : amplitude of tone a); 4th plot : amplitude of tone b); 5th plot : amplitude of tone c); Bottom plot : residues. Red : tone a) estimate; Green : tone b) estimate; Blue : tone c) estimate; Black : DFT.	78
4.23	JET pulse 68558 - Three tone adaptive rEKFFT - Top plot : raw signal spectrogram (\log_{10} scale); 2nd plot : estimated frequencies; 3rd plot : amplitude of tone a); 4th plot : amplitude of tone b); 5th plot : amplitude of tone c); Bottom plot : residues. Red : tone a) estimate; Green : tone b) estimate; Blue : tone c) estimate; Black : DFT.	78
4.24	JET pulse 68558 - Three tone adaptive UKFFT - Top plot : raw signal spectrogram (\log_{10} scale); 2nd plot : estimated frequencies; 3rd plot : amplitude of tone a); 4th plot : amplitude of tone b); 5th plot : amplitude of tone c); Bottom plot : residues. Red : tone a) estimate; Green : tone b) estimate; Blue : tone c) estimate; Black : DFT.	79
4.25	Example of a $q = m/n = 3/2$ flux surface and field line	82
4.26	Data flow diagram for the KF based mode number analysis algorithm . . .	87

4.27	Noise pdf - experimental measurements (solid) and fitted gaussian (dashed) a) T002 coil located at $\phi = 42.15^\circ$; b) H305 coil located at $\phi = 110.38^\circ$; c) T007 coil located at $\phi = 222.17^\circ$; d) T009 coil located at $\phi = 290.37^\circ$	88
4.28	Spatio-temporal data unfolding mechanism - a) signal phase as a function of coil location for 4 consecutive instants; b) measured signal as a function of coil location for 4 consecutive instants.	89
4.29	Degenerate modes of different modulated amplitudes observed at the same modulated frequency - a) Modulated frequency at which the modes are observed; b) Target (dashed) and estimated (solid) amplitudes for each mode.	91
4.30	Absolute value of the mean amplitude estimation error as a function of the process covariance parameter χ	92
4.31	Observability of the system determines the estimation performance - a) worst case amplitude estimation error; b) absolute value of the observabil- ity matrix's determinant.	94
4.32	Toroidal mode number analysis - a) STDFT of a pick-up coil (\log_{10} scale); b) Mode analysis using the coherence spectrum; c) Frequency estimation given by the STDFT (black) and by the EKFFT (blue); d) Amplitude estimation given by the STDFT (black) and KFMP amplitude estimation of the $n=0$ (blue), $n=1$ (green), $n=2$ (red) and $n=3$ (cyan); e) EKFFT estimated residue for coil at toroidal position ϕ_1 (blue), ϕ_2 (green), ϕ_3 (red) and ϕ_4 (cyan).	95
4.33	Toroidal mode number analysis - a) STDFT of a pick-up coil (\log_{10} scale); b) Mode analysis using the coherence spectrum; c) Frequency estimation given by the STDFT (black) and by the EKFFT (blue); d) Amplitude estimation given by the STDFT (black) and KFMP amplitude estimation of the $n=0$ (blue), $n=1$ (green), $n=2$ (red) and $n=3$ (cyan); e) EKFFT estimated residue for coil at toroidal position ϕ_1 (blue), ϕ_2 (green), ϕ_3 (red) and ϕ_4 (cyan).	96
4.34	Toroidal mode number analysis - a) STDFT of a pick-up coil (\log_{10} scale); b) Mode analysis using the coherence spectrum; c) Frequency estimation given by the STDFT (black) and by the EKFFT (blue); d) Amplitude estimation given by the STDFT (black) and KFMP amplitude estimation of the $n=0$ (blue), $n=1$ (green), $n=2$ (red) and $n=3$ (cyan); e) EKFFT estimated residue for coil at toroidal position ϕ_1 (blue), ϕ_2 (green), ϕ_3 (red) and ϕ_4 (cyan).	97
5.1	The error field correction coils (represented in red) at JET. <i>Courtesy of EFDA-JET</i>	100
5.2	EFCC current controller in the JET context.	101
5.3	EFCC current controller Level-1 interface.	102
5.4	(a) Histogram of the application's measured cycle time and (b) Probability of having a processing cycle with an absolute jitter larger than a given percentage of the application's nominal cycle time.	104
5.5	Time stamping and synchronisation mechanism.	105
5.6	Control algorithm simplified data flow diagram.	106
5.7	Time trace - (blue) Current measurement and (red) KF current estimate.	107

5.8	KF frequency response - (a) Normalized amplitude error and (b) Delay normalized to the period.	108
5.9	Saddle coils used in each of the 4 equally spaced octants to calculate the mode-lock signal. <i>Courtesy of EFDA-JET</i>	109
5.10	Mode-lock compensation (color legend in the figure) - continuous line represents the raw signals and dots represents the compensated signals.	110
5.11	Automatic current reference adjustment based on mode-lock compensation.	111
5.12	Slew rate - (blue) estimated current; (green) adjusted reference and (red) slope limited reference	113
5.13	Pre-configured voltage control - (a) Readback of the voltage request acknowledged by the VCVS and (b) Current in the EFCC circuit.	116
5.14	VCVS linear model coefficients.	117
5.15	VCVS linear operational domain - red indexes are labels to identify the regions in the frontier of the operational domain (1 – 8) and inside the operational domain (0).	117
5.16	Commissioning of the reference multiplexing logic.	119
5.17	Pre-configured current reference.	120
5.18	VCVS model assessment - (blue) VCVS output; (green) Controller request; (red) Predicted VCVS output - (a) MPC with $\mu = \xi = 10^{-4}$, (b) MPC with $\mu = \xi = 10^{-3}$, (c) MPC with $\mu = \xi = 10^{-2}$, (d) PID with $K_p = 5, K_i = 20, K_d = 0$	121
5.19	VCVS model assessment - (blue) MPC with $\mu = \xi = 10^{-4}$, (green) MPC with $\mu = \xi = 10^{-3}$, (red) MPC with $\mu = \xi = 10^{-2}$, (cyan) PID with $K_p = 5, K_i = 20, K_d = 0$ - (a) Normalized amplitude error; (b) Delay normalized to period	122
5.20	Control performance - (blue) MPC with $\mu = \xi = 10^{-4}$, (green) MPC with $\mu = \xi = 10^{-3}$, (red) MPC with $\mu = \xi = 10^{-2}$, (cyan) PID with $K_p = 5, K_i = 20, K_d = 0$ and (black) reference - (a) Estimated current; (b) Requested voltage; (c) Cost function; (d) VCVS operational space index (refer to figure 5.15).	123
5.21	Control performance - (blue) MPC with $\mu = \xi = 10^{-4}$, (green) MPC with $\mu = \xi = 10^{-3}$, (red) MPC with $\mu = \xi = 10^{-2}$, (cyan) PID with $K_p = 5, K_i = 20, K_d = 0$ - (a) Normalized amplitude error; (b) Delay normalized to period	123

List of Acronyms

Amplitude Modulation (AM)
Avalanche Photo-Diode (APD)

Discrete Fourier Transform (DFT)
Dual Phase Lock-in (DPL)

Electron Cyclotron Current Drive (ECCD)
Electron Cyclotron Resonance Heating (ECRH)
Error Field Correction Coils (EFCC)
Extended Kalman Filter (EKF)
Extended Kalman Filter Frequency Tracker (EKFFT)
Edge Localized Modes (ELM)

Fast Fourier Transform (FFT)
Finite Impulse Response (FIR)
Frequency Tracker (FT)
Full Width at Half Maximum (FWHM)

iterated Extended Kalman Filter (iEKF)
Infinite Impulse Response (IIR)

Joint European Torus (JET)

Kalman Filter (KF)
Kalman Filter Harmonic Estimator (KFHE)

Low-Pass Filter (LPF)

Maximum A Posteriori (MAP)
Multithreaded Application Real-Time executor (MARTE)
Minimum Mean Square Error (MMSE)
Model Predictive Control (MPC)
Motional Stark Effect (MSE)

Photo-Elastic Modulators (PEMs)
Proportional Integral Derivative (PID)
Pulse Termination Network (PTN)

robust Extended Kalman Filter (rEKF)
robust Extended Kalman Filter Frequency Tracker (rEKFFT)
Radio Frequency (RF)
Resonant Field Amplification (RFA)
Resonant Magnetic Perturbations (RMPs)
Real-Time (RT)
Real-Time Central Controller (RTCC)
Real-Time Data Network (RTDN)
Real-Time Protection Sequencer (RTPS)

Spatial Discrete Fourier Transform (SDFT)
Session Leader (SL)
Signal to Noise Ratio (SNR)
Single Phase Lock-in (SPL)
Short-Time Discrete Fourier Transform (STDFT)

Unscented Kalman Filter (UKF)
Unscented Kalman Filter Frequency Tracker (UKFFT)

Voltage Controlled Oscillator (VCO)
Voltage Controlled Voltage Source (VCVS)
Versa Module European (VME)

Chapter 1

Introduction

1.1 Nuclear fusion power

Massive political, economical, scientific and engineering effort is currently being put into demonstrating that nuclear fusion power can be harnessed on Earth and be considered as a viable energy source for humanity. This effort is revealed on the International Thermonuclear Experimental Reactor, currently being built in France, an ultimate experimental facility that will feature many of the engineering concepts and components that are relevant for a fusion power-plant and confirm the physics understanding of a burning plasma developed over the last 50 years. The reaction that currently congregates most consensus amongst the scientific community to be used for this purpose is given by equation (1.1). Basically two atomic nuclei (the hydrogen isotopes: deuterium and tritium) collide and undergo a fusion reaction giving rise to a single heavier nucleus (Helium) plus a neutron. In this particular reaction, because the mass of the products is lighter than the sum of the mass of the reactants (energy difference accounted for by the different binding energies of the reactants and the products) and due to Einstein's celebrated $E = mc^2$, 17.6 MeV are released in the form of kinetic energy.



In order to achieve fusion, the electrostatic force repelling the nuclei must be overcome so that fusion can occur. This is possible if the nuclei have sufficient kinetic energy to overcome the Coulomb barrier in the event of a collision. Clearly, since the repulsive force is directly proportional to the product of the charge of the nuclei, the lower the atomic number the less the kinetic energy required. One method of increasing the reactants' kinetic energy is by heating them. The rate at which fusion reactions occur is directly proportional to the reactant's density (number of particles per unit volume) and to the collision's effective cross-section. Figure 1.1 illustrates the benefit of the D-T reaction in terms of the effective collision cross-section as a function of the nuclei's energy when compared to other candidate reactions. A thorough coverage and additional bibliography on the rate of fusion reactions can be found in [Wess 11]. It is important to note that, at temperatures of tens or hundreds of keV, atoms become fully ionised thus leaving the electrons unbound from the nuclei in an ionized gas commonly known as a plasma [Chen 84].

Regarding resources, the mass of deuterium in the oceans is estimated to be $\approx 4 \times 10^{16}$ kg [Wess 11] thus vastly abundant and also cheap to extract. Tritium is not available in nature because it undergoes a radioactive decay with a 12.3 year half-life. For this

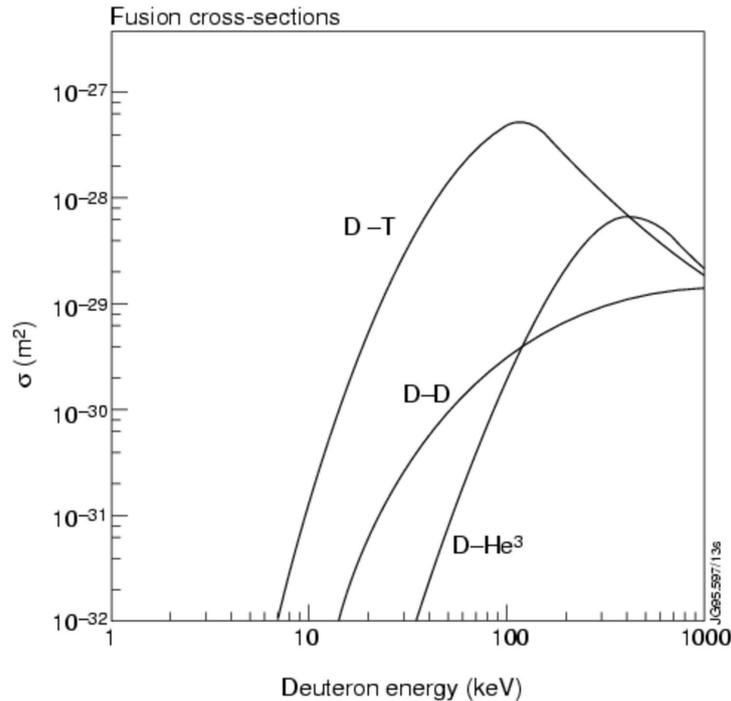


Figure 1.1: Effective collision cross-section of different types of fusion reactions. *Courtesy of EFDA-JET.*

reason, most designs of fusion reactors include a vacuum vessel blanket of a lithium-rich material (for which there are large reserves at competitive prices) in order to breed tritium by using the fast neutrons produced in fusion reactions [Wess 11].

1.2 Tokamaks

Whereas the burning plasma in the Sun's core is self-contained on account of the Sun's massive gravitational field, at the required temperatures of 100 million degrees centigrade, approximately four to ten times higher than the temperature in the core of the sun, the use of material walls to confine the plasma is not an option on earth, even at typical operating densities of $10^{19} - 10^{20} \text{ m}^{-3}$. Magnetic-based confinement has demonstrated to be a viable option. Because a plasma is a fully ionised medium, charged particles undergo Larmor gyration, see figure 1.2(b), around magnetic field lines with a radius directly proportional to their mass and inversely proportional to the field intensity. This mechanism reduces the plasma heat load on containment structures to a range which is manageable by present technological capabilities. Of the possible magnetic confinement configurations explored to date, the tokamak is widely acknowledged as the most promising when it comes to demonstrating the confinement and power generation from a fusion burning plasma.

A tokamak is a torus shaped container, see figures 1.3 and 1.4, where the plasma is confined as a consequence of externally applied and inductively driven magnetic fields (see figure 1.2(a)). In fact, strong external magnetic fields are produced by currents flowing in dedicated coils that are arranged in a, quasi axis-symmetric toroidal configuration. The resultant magnetic field (toroidal magnetic field) is nonetheless insufficient to confine the hot plasma due to the variation of the magnetic field strength with the major radius of the device (distance from the vertical symmetry axis of the device). To contrast the

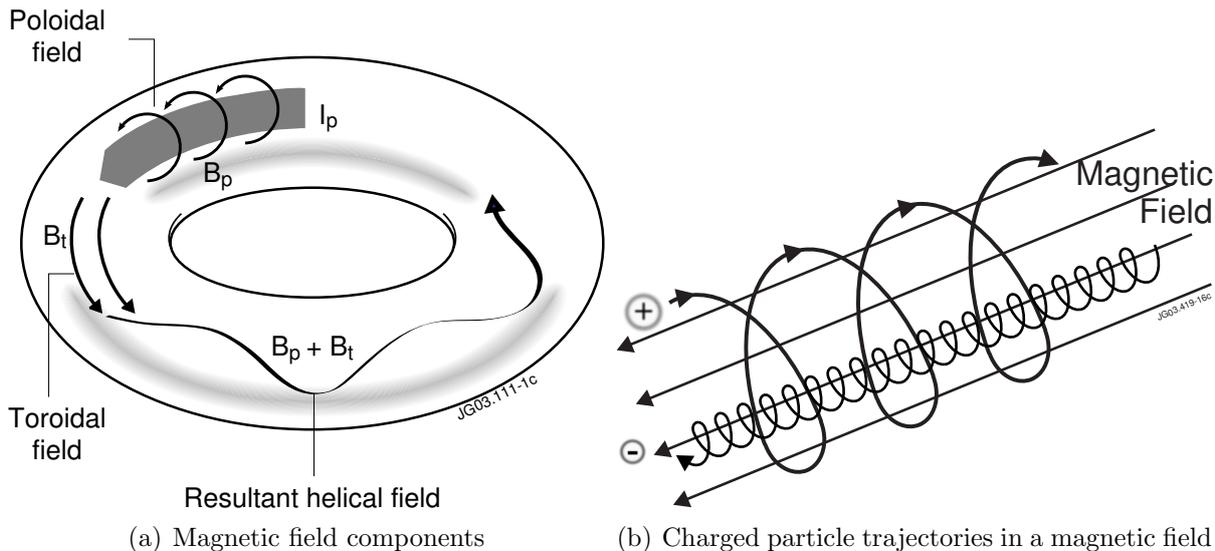


Figure 1.2: Particle confinement. *Courtesy of EFDA-JET.*

resulting outward force exerted on the confined plasma, a *poloidal* magnetic field is deemed necessary to ensure a steady balance between the plasma pressure and the magnetic forces. Such a poloidal component results from the inductive drive of the toroidal current via transformer effect using dedicated *ohmic* coils as the primary circuit and with the plasma acting as the secondary circuit. An illustrative example for an iron core tokamak is shown in figure 1.4, evidencing the most important magnetic circuits including the coils generating the horizontal/vertical magnetic fields used for the positioning and shaping of the plasma.

Ohmic heating, self-generated by the plasma due to its finite electrical conductivity and the toroidal plasma current, is not relevant for plasma temperatures exceeding the keV since the plasma conductivity increases strongly with an increasing plasma temperature. As such, reactor relevant plasmas, for which a strong heating component will be provided by the fusion born alpha particles, must be supplemented by additional heating systems. These include neutral beam injection (NBI) and resonant electromagnetic heating at the ion cyclotron (30-50 MHz), lower hybrid (3-6 GHz) and electron cyclotron (70-170 GHz) frequencies.

The inner poloidal field coils, as previously suggested, are used to control the plasma current. The remaining poloidal field coils are current rings concentric with the torus and are used to position and shape the plasma.

Tokamaks, invented in the former Soviet Union in the 1950s, have diffused around the world in many sizes and shapes. The larger, and one of the few with tritium fuelling capabilities, is the Joint European Torus (JET) [Roma 13] located at Culham, United Kingdom.

1.3 Real-time systems in the world's largest tokamak

Experimental machines as complex as tokamaks require, and rely, on real-time systems for multiple purposes including diagnostics, control, protection or physics processing nodes.

As far as diagnostics are concerned, they provide the measurement support for the feed-back control of the experiment while simultaneously enabling post-pulse physics studies.

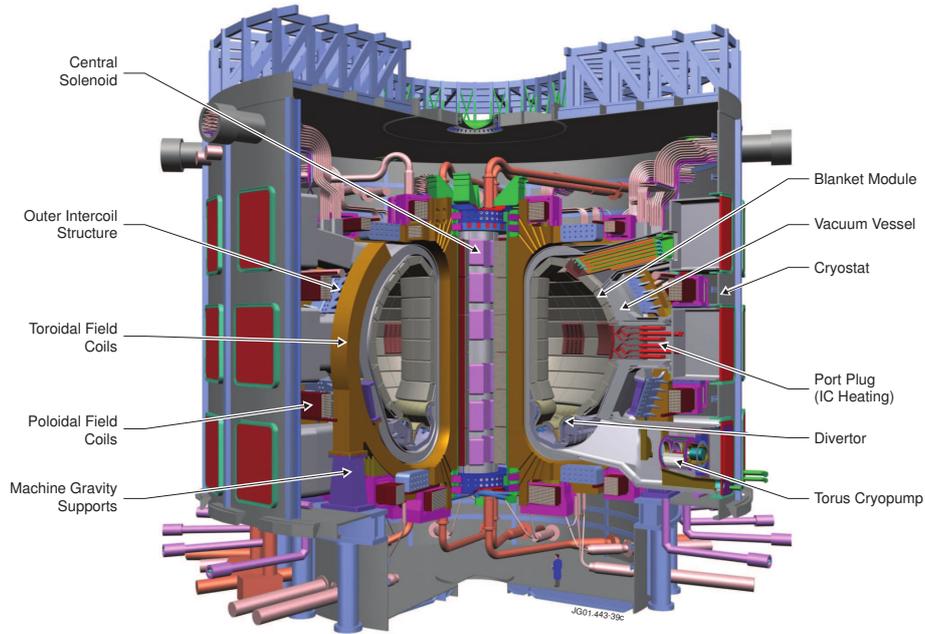


Figure 1.3: Tokamak sketch. *Courtesy of EFDA-JET.*

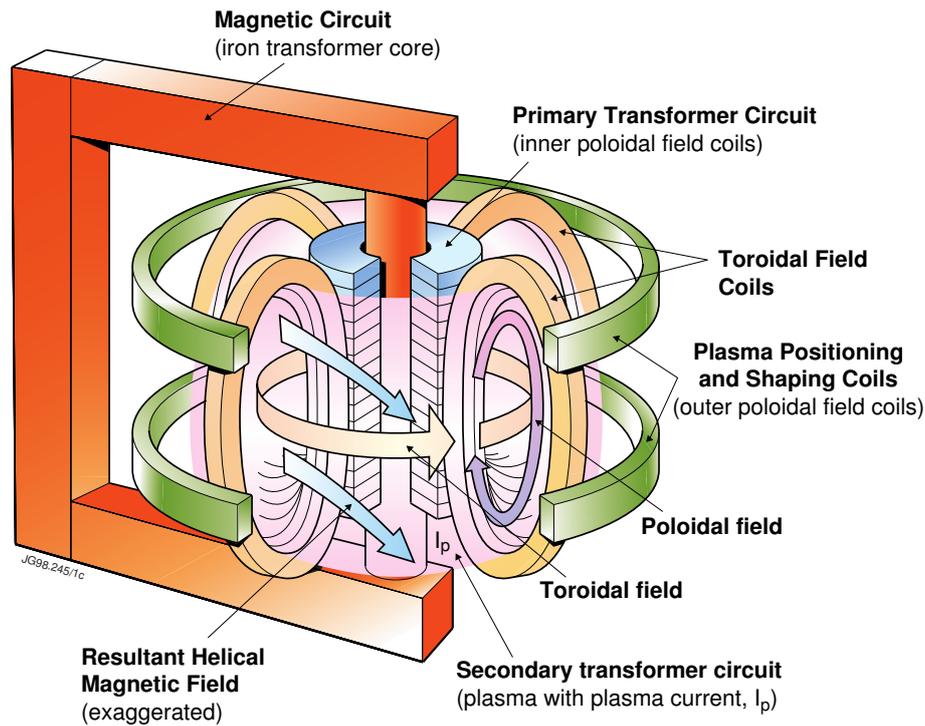


Figure 1.4: Tokamak transformer. *Courtesy of EFDA-JET.*

Magnetic and density measurements play the most fundamental role in terms of machine operations. Magnetics, for instance, are essential for the control of the plasma current, position and shape [Begh 05, Ambr 05, Ario 05]. Density measurements are typically provided by the interferometer and the Thomson scattering diagnostics but can also be obtained from, e.g., the reflectometry or the lithium beam diagnostics. Accurate and responsive density control is of fundamental importance for tokamak operation, affecting operational aspects like the coupling of heating power, stability limit avoidance, fusion

gain optimization (e.g. operation close to Greenwald limit), protection of shine-through regions in machines with neutral beam heating and the scientific efficacy and reproducibility. Others, such as the Motional Stark Effect (MSE) [Alve 04] or the polarimeter may also assist in the real-time determination of the safety factor¹ profile [Zabe 02] in a combined real-time physics processing scheme performed in dedicated node(s).

When it comes to control, the aforementioned magnetic control of the plasma current, position and shape is certainly essential however, in plasmas with a non-circular cross section, a configuration known to be vertically unstable [Shaf 63, Zakh 73], the control of the plasma's vertical position is also crucial [Lenn 97]. In addition, the other main control in tokamak experiments is the control of the additional heating, crucial in order to achieve the experiment's scientific goals and, simultaneously, to protect the machine in case of abnormal events. Feedback control in tokamaks can be categorised within three main scopes: operational, protection and experimental. Often the first two coincide. The first type of control is in general connected with the control of basic quantities, such as plasma current and toroidal field, required for any experiment. Protective actions are essentially related to the mechanical integrity of the machine, for instance, reacting in the event of a short-circuit in one of the toroidal field coils creating large asymmetrical forces or to alarms warning about an imminent disruption [Dorm 13]. Recently at JET, with the advent of the installation of the ITER-like wall, the increased concern of thermal loads on plasma facing components led to the development of an entire new set of real-time systems [Arno 12, Step 11, Alve 12a, Jouv 11, Alve 12b]. From the experimental point of view, real-time control can be used to assist in physics studies by controlling, for instance, the spatial profile of relevant physical quantities [Joff 03, Felt 05], Resonant Magnetic Perturbations (RMPs) [Alve 11b] or the control of the ratio between the plasma pressure and magnetic forces [Joff 03].

The physics real-time processing nodes are basically the ones that enable the experimental type of control. At JET they can provide, for instance, the full equilibrium reconstruction (taking the plasma, vacuum and conduction structures into account) in real-time [Blum 08, Mazo 10], the safety factor spatial profile [Zabe 02] or the plasma's internal inductance and stored energy [Bara 04].

A common denominator across all the aforementioned types of real-time systems is the requirement for accurate, resilient and low latency real-time calculations performed at adequate output rates. In terms of resilience it is worth noting that, for example Edge Localized Modes (ELM) which are common in tokamak-based nuclear fusion research, often perturb a whole suite of diagnostic measurements. This instability causes a quasi-periodic relaxation of (mainly temperature and density) profiles at the plasma edge due to the steep gradients and edge current density formed in high confinement plasma regimes. The relaxation event (also known as ELM crash) occurs in a much faster time scale (typically tens of microseconds) than the inter-crash period (typically tens of milliseconds). At the time of the crash, some of the edge plasma particles and stored energy are lost to the periphery and eventually reach the plasma facing components, predominantly in the divertor region. The resultant increase in radiation can impact severely on radiation sensitive diagnostics such as the MSE diagnostic. Without entering into fine detail, after the ELM crash, it is possible that some stray radiation from the divertor becomes polarized once reflected in metallic plasma facing components and contaminates strongly the polarized light components to which the MSE is sensitive. This can greatly impact the pitch angle estimations derived with the diagnostic. Real-time signal processing tools

¹A measurement of the magnetic field line bending which is strongly connected with plasma stability.

capable of ameliorating the effect of these disturbances are certainly desirable.

1.4 Real-time signal processing

Real-time signal processing is a common requirement in tokamak fusion research. It is required by: (a) diagnostics providing analog, digital or network-based real-time measurements; (b) physics processing nodes combining multiple diagnostic measurements to infer derived physical quantities; (c) controllers using results from the previous systems, eventually conditioning them and subsequently performing their model calculations and actuator requests.

Signal processing techniques are used to achieve standard objectives such as filtering, demodulation, detection, spectral estimation and time-frequency analysis. Depending on the purpose, constraints and context, choosing the appropriate set of signal processing techniques is crucial for achieving the best results. Statistical signal processing, for instance, is a field comprising techniques based on stochastic process modelling to analyze signal behaviour. Signals are modelled as linear or non-linear combinations of time series representing deterministic and stochastic components. Relevant information is extracted using techniques which incorporate knowledge inferred from the analysis of signals' statistical properties.

In the particular context of this work the dominant scope will be filtering, i.e., extracting information about a (potentially multi-dimensional) quantity x at a given time t , given a set of measurements z (with some degree of dependence/correlation with x) up to and including t . Considering a system whose state can be fully characterized, at a given time index k , by state vector \mathbf{x}_k and for which a set of system measurements $\mathbf{y}_{0...k}$ are available up to k , then the goal is to find the optimal state estimate $\mathbf{x}_k = \mathbf{x}_k^{optim}$ with respect to a predefined criterion. Clearly some theoretical knowledge on how the measurements are related to the system state is required. However, if in addition, some theoretical knowledge is available on how the system state evolves from one time instant to the next, both contributions can be combined to provide more accurate estimations than the ones that would have been obtained with just measurement or state transition individual knowledge. In case the state transition model and measurement-to-state *mapping* are linear and in the presence of *white disturbances*², a linear tractable solution (the Kalman Filter (KF)) is derivable. However, in the general case where some/all of the previous premises are not satisfied, such solutions are generally not tractable.

This dissertation presents applications developed, using KF-based methods, to meet signal processing requirements in tokamak fusion research and is organized as follows: in chapter 2 the KF is thoroughly described and derived both using the Minimum Mean Square Error (MMSE) criterion and in the context of Bayesian analysis. Chapter 3 essentially describes the Kalman Filter Harmonic Estimator (KFHE) and presents some applications based on it for real-time signal processing of the MSE diagnostic at JET (section 3.3.1), the coherence analysis of magnetic fluctuations at ISTTOK (section 3.3.2) as well as synchronous and resonance detection for the toroidal Alfvén eigenmode (TAE) antenna diagnostic at JET (section 3.3.3). Chapter 4 introduces an implementation of a non-linear KF-based Frequency Tracker (FT). In section 4.3 studies are presented on the behaviour of the Extended Kalman Filter Frequency Tracker (EKFFT) in the context of

²Disturbances following a distribution of zero-mean and serially uncorrelated random variables with uniform power spectral density.

simultaneous multi-tone estimation. Section 4.4.1 deals with a non-linear approach for the TAE signal processing previously addressed in section 3.3.3 and sections 4.4.2 and 4.4.3 deal, respectively, with real-time frequency tracking and mode number analysis of MagnetoHydroDynamic (MHD) activity in JET plasmas. In chapter 5, the real-time controller of RMPs, employing a KF-based observer, is described and finally, in chapter 6, a brief discussion of the results presented in this dissertation and conclusions are presented.

Chapter 2

Kalman-based state estimation

This chapter provides the theoretical basis for the various KF flavours whose implementations are described in subsequent chapters. Section 2.1 introduces the filtering problem and section 2.2 describes the KF comprehensively including two independent derivations. Finally, in section 2.3, a set of non-linear extensions of the KF is presented along with a comprehensive derivation of their pillar, the Extended Kalman Filter (EKF).

2.1 Brief introduction to the filtering problem

Typical usage of a filter addresses the problem of the simultaneous estimation of a set of quantities that characterize a system at a given time instant, given the history of known inputs and noise polluted observations. This set is called the system's *state*. Information on the system dynamics, i.e., the temporal evolution of the system state and its connection with the observations, is ideally available and incorporated into the estimation process. The estimates obtained via this process usually conform to some optimal criteria that typically minimizes the estimation error in some sense. The MMSE, for instance. In the Bayesian formalism [Cand 09], the filter evolves the probability density function (pdf) of the system state conditioned to the available measurements. In the most generic case where the system dynamics is non-linear, the pdf cannot be parameterized by a finite amount of moments and therefore needs to be fully propagated in order to preserve the information about the system. Under very special circumstances though, the KF hypothesis being one of them, not only the parameterization of the pdf is possible but also the propagation of these parameters obey well known relations. In these cases, accurate estimations are possible without the unpractical propagation of the entire pdf.

Denoting the discrete time variable by k , the system state by $\boldsymbol{\Omega}_k$, the set of system inputs by \mathbf{u}_k , the set of system measurements by \mathbf{z}_k , the process noise by \mathbf{w}_k and the measurement noise by \mathbf{v}_k , it is common to write the discretized form of both the state transition equation, i.e., the equation which evolves the system state from one sample to the next, and the measurement equation, i.e., the equation that relates the system state to measurable quantities, respectively, as equations (2.1) and (2.2). It is worth noting that \mathbf{w}_k is a stochastic term that represents the uncertainty in the knowledge of the state transition dynamics and that \mathbf{v}_k is a term representing simply the ordinary noise in system measurements.

$$\boldsymbol{\Omega}_k = \mathbf{f}_{k-1}(\boldsymbol{\Omega}_{k-1}, \mathbf{u}_{k-1}, \mathbf{w}_{k-1}) \quad (2.1)$$

$$\mathbf{z}_k = \mathbf{h}_k(\boldsymbol{\Omega}_k, \mathbf{v}_k) \quad (2.2)$$

Implicitly assumed in equation (2.1) is the modelling of the system's state transition as a Markov process meaning that the present state is assumed to depend solely on the previous state and previous inputs rather than their entire past history.

2.2 Kalman filter

The KF is a recursive *predictor-corrector* estimator that incorporates information from both measurements and theoretical models, combining those to provide an optimal estimate of a system's state at a given time instant. The KF, also known as the *linear quadratic estimator*, uses a prior state estimation combined with the theoretical model for the system's state transition to predict its current state in the *prediction* stage. It then revises its own prediction by incorporating information from system measurements in the *corrector* stage. The KF's validity and optimality are nonetheless only guaranteed under certain assumptions depending on the formalism/context.

In this section, it is shown how the KF equations can be derived using two different and independent approaches. One based on linear systems theory and mean square minimization arguments and another within the context of the bayesian formalism and the *maximum a posteriori* criterion. The treatment presented herein will be restricted to the discrete time case mainly due to its immediate applicability to the subsequent practical implementations described in this dissertation.

2.2.1 Minimum mean square error derivation

Let one consider a system whose deterministic state transition model is linear, Markovian¹ and given by equation (2.3) where k denotes the discrete time index, \mathbf{x}_k^{true} is the $N \times 1$ vector of quantities describing the true system state, \mathbf{A}_{k-1}^{true} is the $N \times N$ matrix describing the system's state transition in the absence of exogenous inputs, \mathbf{B}_{k-1}^{true} is the $N \times M$ matrix relating external inputs with the system state and, finally, \mathbf{u}_{k-1} is the $M \times 1$ vector of external inputs. \mathbf{A}_{k-1}^{true} and \mathbf{B}_{k-1}^{true} provide a full, and infinitely accurate, description of the deterministic system dynamics.

$$\mathbf{x}_k^{true} = \mathbf{A}_{k-1}^{true} \mathbf{x}_{k-1}^{true} + \mathbf{B}_{k-1}^{true} \mathbf{u}_{k-1} \quad (2.3)$$

In addition, let one consider a linear relation between the system state and a set of noise polluted observable quantities given by equation (2.4) where \mathbf{z}_k is the $L \times 1$ vector containing the set of system measurements, \mathbf{C}_k is the $L \times N$ matrix relating the system's state with the measurements and \mathbf{v}_k is the $L \times 1$ set of noise terms contaminating the measurements.

$$\mathbf{z}_k = \mathbf{C}_k \mathbf{x}_k^{true} + \mathbf{v}_k \quad (2.4)$$

Having a linear Markovian system described by (2.3) and (2.4), let one now define an *a priori* (often referred to as *prior*) state estimate $\hat{\mathbf{x}}_k$ and an *a posteriori* (often referred to as *posterior*) state estimate \mathbf{x}_k . The former, given by equation (2.5), is the state estimate

¹A Markov process is, by definition, a process by which, given the knowledge of the present state, the future state of a system is independent of all other previous states.

before the set of measurements \mathbf{z}_k is available and the latter is the state estimate after \mathbf{z}_k becomes available. Given the posterior state estimate \mathbf{x}_{k-1} for $k-1$, and making use of the available theoretical knowledge of the system dynamics, $\hat{\mathbf{x}}_k$ represents a prediction of the system state at k before measurements \mathbf{z}_k are available. In this equation, \mathbf{w}_{k-1} is an $N \times 1$ noise vector representative of the uncertainty in the accuracy of the model and \mathbf{A}_{k-1} and \mathbf{B}_{k-1} represent the best available theoretical description of the deterministic system dynamics.

$$\hat{\mathbf{x}}_k = \mathbf{A}_{k-1}\mathbf{x}_{k-1} + \mathbf{B}_{k-1}\mathbf{u}_{k-1} + \mathbf{w}_{k-1} \quad (2.5)$$

Let one now postulate that the posterior state estimate \mathbf{x}_k is a linear combination of the prediction (prior state estimate) $\hat{\mathbf{x}}_k$ and the measurements \mathbf{z}_k given by equation (2.6).

$$\mathbf{x}_k = \mathbf{K}'_k \hat{\mathbf{x}}_k + \mathbf{K}_k \mathbf{z}_k \quad (2.6)$$

Substituting (2.4) in the previous yields (2.7).

$$\mathbf{x}_k = \mathbf{K}'_k \hat{\mathbf{x}}_k + \mathbf{K}_k \mathbf{C}_k \mathbf{x}_k^{true} + \mathbf{K}_k \mathbf{v}_k \quad (2.7)$$

Defining the posterior state estimation error $\boldsymbol{\epsilon}_k$ and using the previous yields (2.8).

$$\begin{aligned} \boldsymbol{\epsilon}_k &= \mathbf{x}_k^{true} - \mathbf{x}_k \\ &= \mathbf{x}_k^{true} - \mathbf{K}'_k \hat{\mathbf{x}}_k - \mathbf{K}_k \mathbf{C}_k \mathbf{x}_k^{true} - \mathbf{K}_k \mathbf{v}_k \\ &= (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \mathbf{x}_k^{true} - \mathbf{K}'_k \hat{\mathbf{x}}_k - \mathbf{K}_k \mathbf{v}_k \end{aligned} \quad (2.8)$$

Letting $\mathbf{K}'_k = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)$ allows one to write both (2.9), which intuitively says that the posterior state estimate is equal to the prior state estimate corrected by a quantity proportional to the difference between the predicted measurements $\hat{\mathbf{z}}_k = \mathbf{C}_k \hat{\mathbf{x}}_k$ and the actual set of measurements \mathbf{z}_k , and (2.10).

$$\mathbf{x}_k = \hat{\mathbf{x}}_k + \mathbf{K}_k (\mathbf{z}_k - \mathbf{C}_k \hat{\mathbf{x}}_k) \quad (2.9)$$

$$\boldsymbol{\epsilon}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) (\mathbf{x}_k^{true} - \hat{\mathbf{x}}_k) - \mathbf{K}_k \mathbf{v}_k \quad (2.10)$$

Defining the prior state estimation error $\hat{\boldsymbol{\epsilon}}_k = \mathbf{x}_k^{true} - \hat{\mathbf{x}}_k$ allows one to write the previous as (2.11).

$$\boldsymbol{\epsilon}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\boldsymbol{\epsilon}}_k - \mathbf{K}_k \mathbf{v}_k \quad (2.11)$$

This defines a relation between the posterior and the prior state estimation errors. Taking the expected value of the previous and assuming that \mathbf{v}_k is a random variable drawable from a zero-mean joint distribution gives (2.12).

$$E[\boldsymbol{\epsilon}_k] = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) E[\hat{\boldsymbol{\epsilon}}_k] \quad (2.12)$$

The previous implies that the posterior state estimation is unbiased, i.e. $E[\boldsymbol{\epsilon}_k] = \mathbf{0}$, independently of \mathbf{K}_k , if and only if $E[\hat{\boldsymbol{\epsilon}}_k] = \mathbf{0}$.

Let one now expand $\hat{\boldsymbol{\epsilon}}_k = \mathbf{x}_k^{true} - \hat{\mathbf{x}}_k$, by substituting (2.3) and (2.5), to write (2.13).

$$\hat{\boldsymbol{\epsilon}}_k = \mathbf{A}_{k-1}^{true} \mathbf{x}_{k-1}^{true} + \mathbf{B}_{k-1}^{true} \mathbf{u}_{k-1} - \mathbf{A}_{k-1} \mathbf{x}_{k-1} - \mathbf{B}_{k-1} \mathbf{u}_{k-1} - \mathbf{w}_{k-1} \quad (2.13)$$

Assuming the average system dynamics model is accurate, i.e. $\mathbf{A}_{k-1} = \mathbf{A}_{k-1}^{true}$ and $\mathbf{B}_{k-1} = \mathbf{B}_{k-1}^{true}$, and noting that $\boldsymbol{\epsilon}_{k-1} = \mathbf{x}_{k-1}^{true} - \mathbf{x}_{k-1}$, the previous can be written as (2.14).

$$\hat{\boldsymbol{\epsilon}}_k = \mathbf{A}_{k-1}^{true} \boldsymbol{\epsilon}_{k-1} - \mathbf{w}_{k-1} \quad (2.14)$$

Again, taking the expected value of the previous and assuming that \mathbf{w}_{k-1} is a random variable drawable from a zero-mean joint distribution yields (2.15).

$$E[\hat{\boldsymbol{\epsilon}}_k] = \mathbf{A}_{k-1}^{true} E[\boldsymbol{\epsilon}_{k-1}] \quad (2.15)$$

Therefore, replacing the previous in (2.12) yields (2.16).

$$E[\boldsymbol{\epsilon}_k] = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \mathbf{A}_{k-1}^{true} E[\boldsymbol{\epsilon}_{k-1}] \quad (2.16)$$

So, if \mathbf{v}_k and \mathbf{w}_k are drawable from zero-mean joint distributions and the average system model used for the prior state estimate is accurate, should $E[\boldsymbol{\epsilon}_{k-1}] = \mathbf{0}$ for some $k-1$, then $E[\boldsymbol{\epsilon}_k] = \mathbf{0} \quad \forall k$, independently of \mathbf{K}_k .

An unbiased estimator is certainly desirable but an even more desirable estimator is one for which the variance of the state estimation error is as small as possible. In order to derive such an estimator, let one start by writing the covariance matrix of the posterior state estimation error defined by $\mathbf{P}_k = E[\boldsymbol{\epsilon}_k \boldsymbol{\epsilon}_k^T]$, where $\boldsymbol{\epsilon}_k^T$ denotes the transpose of $\boldsymbol{\epsilon}_k$. Then, using (2.11) one can write (2.17).

$$\begin{aligned} \mathbf{P}_k &= E \left[[(\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\boldsymbol{\epsilon}}_k - \mathbf{K}_k \mathbf{v}_k] [(\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\boldsymbol{\epsilon}}_k - \mathbf{K}_k \mathbf{v}_k]^T \right] \\ &= E \left[[(\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\boldsymbol{\epsilon}}_k - \mathbf{K}_k \mathbf{v}_k] \left[\hat{\boldsymbol{\epsilon}}_k^T (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)^T - \mathbf{v}_k^T \mathbf{K}_k^T \right] \right] \\ &= E \left[(\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\boldsymbol{\epsilon}}_k \hat{\boldsymbol{\epsilon}}_k^T (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)^T - (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\boldsymbol{\epsilon}}_k \mathbf{v}_k^T \mathbf{K}_k^T - \right. \\ &\quad \left. \mathbf{K}_k \mathbf{v}_k \hat{\boldsymbol{\epsilon}}_k^T (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)^T + \mathbf{K}_k \mathbf{v}_k \mathbf{v}_k^T \mathbf{K}_k^T \right] \\ &= (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) E[\hat{\boldsymbol{\epsilon}}_k \hat{\boldsymbol{\epsilon}}_k^T] (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)^T - (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) E[\hat{\boldsymbol{\epsilon}}_k \mathbf{v}_k^T] \mathbf{K}_k^T - \\ &\quad \mathbf{K}_k E[\mathbf{v}_k \hat{\boldsymbol{\epsilon}}_k^T] (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)^T + \mathbf{K}_k E[\mathbf{v}_k \mathbf{v}_k^T] \mathbf{K}_k^T \end{aligned} \quad (2.17)$$

One had previously assumed \mathbf{v}_k to be zero-mean. In addition, let one also assume that \mathbf{v}_k is white, i.e., its sequential realizations are uncorrelated. This can be mathematically expressed by (2.18) where i and j are sequence (time) indexes, δ_{ij} is the Kroenecker delta and \mathbf{R}_i is the measurement noise covariance matrix.

$$E[\mathbf{v}_i \mathbf{v}_j^T] = \delta_{ij} \mathbf{R}_i \quad (2.18)$$

The previous, together with the observation that $\hat{\boldsymbol{\epsilon}}_k$ and \mathbf{v}_k are independent implies that $E[\hat{\boldsymbol{\epsilon}}_k \mathbf{v}_k^T] = E[\hat{\boldsymbol{\epsilon}}_k] E[\mathbf{v}_k^T] = \mathbf{0}$. The same arguments hold for $E[\mathbf{v}_k \hat{\boldsymbol{\epsilon}}_k^T]$. If in addition one notes that $E[\hat{\boldsymbol{\epsilon}}_k \hat{\boldsymbol{\epsilon}}_k^T]$ is, by definition, $\hat{\mathbf{P}}_k$, i.e., the prior state estimation error's covariance matrix, and that $E[\mathbf{v}_k \mathbf{v}_k^T] = \mathbf{R}_k$, (2.17) becomes (2.19).

$$\mathbf{P}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\mathbf{P}}_k (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)^T + \mathbf{K}_k \mathbf{R}_k \mathbf{K}_k^T \quad (2.19)$$

Let one now define the quadratic sum of the estimation errors given by (2.20).

$$\|\mathbf{x}_k^{true} - \mathbf{x}_k\|^2 = \sum_{i=1}^N (x_{i,k}^{true} - x_{i,k})^2 = \sum_{i=1}^N \epsilon_{i,k}^2 = \boldsymbol{\epsilon}_k^T \boldsymbol{\epsilon}_k = Tr(\boldsymbol{\epsilon}_k \boldsymbol{\epsilon}_k^T) \quad (2.20)$$

Noting that $E[Tr(\boldsymbol{\epsilon}_k \boldsymbol{\epsilon}_k^T)] = Tr(E[\boldsymbol{\epsilon}_k \boldsymbol{\epsilon}_k^T]) = Tr(\mathbf{P}_k)$, the mean square error cost function J_k is defined by (2.21).

$$J_k = E[\|\mathbf{x}_k^{true} - \mathbf{x}_k\|^2] = E\left[\sum_{i=1}^N (x_{i,k}^{true} - x_{i,k})^2\right] = \sum_{i=1}^N E[(x_{i,k}^{true} - x_{i,k})^2] = Tr(\mathbf{P}_k) \quad (2.21)$$

It all comes down to the unconstrained minimization of a sum of squares which is a convex problem without local or global maxima. One can then find the value of \mathbf{K}_k that minimizes J_k by differentiating it, setting it equal to zero and solving for \mathbf{K}_k .

$$\begin{aligned} \frac{\partial J_k}{\partial \mathbf{K}_k} &= \frac{\partial Tr(\mathbf{P}_k)}{\partial \mathbf{K}_k} \\ &= \frac{\partial Tr[(\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\mathbf{P}}_k (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)^T + \mathbf{K}_k \mathbf{R}_k \mathbf{K}_k^T]}{\partial \mathbf{K}_k} \\ &= \frac{\partial Tr[(\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\mathbf{P}}_k (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)^T]}{\partial \mathbf{K}_k} + \frac{\partial Tr[\mathbf{K}_k \mathbf{R}_k \mathbf{K}_k^T]}{\partial \mathbf{K}_k} \\ &= \frac{\partial Tr[(\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\mathbf{P}}_k (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)^T]}{\partial (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)} \frac{\partial (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k)}{\partial \mathbf{K}_k} + \frac{\partial Tr[\mathbf{K}_k \mathbf{R}_k \mathbf{K}_k^T]}{\partial \mathbf{K}_k} = 0 \end{aligned} \quad (2.22)$$

From matrix calculus one has that $\partial [Tr(ABA^T)] / \partial A = 2AB$ if B is symmetric and $\partial [AB] / \partial A = B^T$. These identities allow one to rewrite (2.22) in the form of (2.23) provided that $\hat{\mathbf{P}}_k$ and \mathbf{R}_k are symmetric.

$$-2(\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \mathbf{P}_{k-1} \mathbf{C}_k^T + 2\mathbf{K}_k \mathbf{R}_k = 0 \quad (2.23)$$

Straightforward algebraic manipulation of the previous finally yields (2.24) for the \mathbf{K}_k that minimizes the state estimate's mean square error.

$$\mathbf{K}_k = \mathbf{P}_{k-1} \mathbf{C}_k^T (\mathbf{C}_k \mathbf{P}_{k-1} \mathbf{C}_k^T + \mathbf{R}_k)^{-1} \quad (2.24)$$

Substituting the previous in (2.19) yields the simplified expression (2.25) for \mathbf{P}_k .

$$\mathbf{P}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\mathbf{P}}_k \quad (2.25)$$

The only thing left to do in order to close the recursive system of expressions that form the KF is to relate $\hat{\mathbf{P}}_k$ with \mathbf{P}_{k-1} . Noting that $\hat{\mathbf{P}}_k = E[\hat{\boldsymbol{\epsilon}}_k \hat{\boldsymbol{\epsilon}}_k^T] = E[(\mathbf{x}_k^{true} - \hat{\mathbf{x}}_k)(\mathbf{x}_k^{true} - \hat{\mathbf{x}}_k)^T]$ and substituting \mathbf{x}_k^{true} by (2.3) along with $\hat{\mathbf{x}}_k$ by (2.5) yields (2.26).

$$\begin{aligned} \hat{\mathbf{P}}_k &= E\left[\left(\mathbf{A}_{k-1} \mathbf{x}_{k-1}^{true} + \mathbf{B}_{k-1}^{true} \mathbf{u}_{k-1} - \mathbf{A}_{k-1} \hat{\mathbf{x}}_{k-1} - \mathbf{B}_{k-1} \mathbf{u}_{k-1} - \mathbf{w}_{k-1}\right) \right. \\ &\quad \left. \left(\mathbf{A}_{k-1} \mathbf{x}_{k-1}^{true} + \mathbf{B}_{k-1}^{true} \mathbf{u}_{k-1} - \mathbf{A}_{k-1} \hat{\mathbf{x}}_{k-1} - \mathbf{B}_{k-1} \mathbf{u}_{k-1} - \mathbf{w}_{k-1}\right)^T\right] \end{aligned} \quad (2.26)$$

Assuming that all unmodelled dynamics are described by the white and zero-mean stochastic term \mathbf{w}_{k-1} , one has that $\mathbf{A}_{k-1} = \mathbf{A}_{k-1}^{true}$ and $\mathbf{B}_{k-1} = \mathbf{B}_{k-1}^{true}$ hence simplifying the previous to (2.27).

$$\hat{\mathbf{P}}_k = E\left[\left[\mathbf{A}_{k-1} (\mathbf{x}_{k-1}^{true} - \hat{\mathbf{x}}_{k-1}) - \mathbf{w}_{k-1}\right] \left[\mathbf{A}_{k-1} (\mathbf{x}_{k-1}^{true} - \hat{\mathbf{x}}_{k-1}) - \mathbf{w}_{k-1}\right]^T\right] \quad (2.27)$$

Noting that $\boldsymbol{\epsilon}_{k-1} = \mathbf{x}_{k-1}^{true} - \widehat{\mathbf{x}}_{k-1}$ and expanding the previous yields (2.28).

$$\begin{aligned}\widehat{\mathbf{P}}_k &= E \left[\mathbf{A}_{k-1} \boldsymbol{\epsilon}_{k-1} \boldsymbol{\epsilon}_{k-1}^T \mathbf{A}_{k-1}^T - \mathbf{A}_{k-1} \boldsymbol{\epsilon}_{k-1} \mathbf{w}_{k-1}^T - \mathbf{w}_{k-1} \boldsymbol{\epsilon}_{k-1}^T \mathbf{A}_{k-1}^T + \mathbf{w}_{k-1} \mathbf{w}_{k-1}^T \right] \\ &= \mathbf{A}_{k-1} E \left[\boldsymbol{\epsilon}_{k-1} \boldsymbol{\epsilon}_{k-1}^T \right] \mathbf{A}_{k-1}^T - \mathbf{A}_{k-1} E \left[\boldsymbol{\epsilon}_{k-1} \mathbf{w}_{k-1}^T \right] - E \left[\mathbf{w}_{k-1} \boldsymbol{\epsilon}_{k-1}^T \right] \mathbf{A}_{k-1}^T + E \left[\mathbf{w}_{k-1} \mathbf{w}_{k-1}^T \right]\end{aligned}\quad (2.28)$$

One had previously assumed \mathbf{w}_k to be zero-mean. In addition, let one also assume that \mathbf{w}_k is white, i.e., its sequential realizations are uncorrelated. This can be mathematically expressed by (2.29) where i and j are sequence (time) indexes, δ_{ij} is the Kroenecker delta and \mathbf{Q}_i is the process's noise covariance matrix.

$$E \left[\mathbf{w}_i \mathbf{w}_j^T \right] = \delta_{ij} \mathbf{Q}_i \quad (2.29)$$

The previous, together with the observation that $\boldsymbol{\epsilon}_{k-1}$ and \mathbf{w}_{k-1} are independent implies that $E \left[\boldsymbol{\epsilon}_{k-1} \mathbf{w}_{k-1}^T \right] = E \left[\boldsymbol{\epsilon}_{k-1} \right] E \left[\mathbf{w}_{k-1}^T \right] = \mathbf{0}$. The same arguments hold for $E \left[\mathbf{w}_{k-1} \boldsymbol{\epsilon}_{k-1}^T \right]$. If in addition one notes that $E \left[\boldsymbol{\epsilon}_{k-1} \boldsymbol{\epsilon}_{k-1}^T \right]$ is, by definition, \mathbf{P}_{k-1} , i.e., the posterior state estimation error's covariance matrix and that $E \left[\mathbf{w}_{k-1} \mathbf{w}_{k-1}^T \right] = \mathbf{Q}_{k-1}$, one can write (2.30).

$$\widehat{\mathbf{P}}_k = \mathbf{A}_{k-1} \mathbf{P}_{k-1} \mathbf{A}_{k-1}^T + \mathbf{Q}_{k-1} \quad (2.30)$$

It is now worth explicitly writing the entire set of equations that form the celebrated Kalman filter. The prediction stage yields:

$$\widehat{\mathbf{x}}_k = \mathbf{A}_{k-1} \mathbf{x}_{k-1} + \mathbf{B}_{k-1} \mathbf{u}_{k-1} \quad (2.31)$$

$$\widehat{\mathbf{P}}_k = \mathbf{A}_{k-1} \mathbf{P}_{k-1} \mathbf{A}_{k-1}^T + \mathbf{Q}_{k-1} \quad (2.32)$$

The Kalman gain is calculated via:

$$\mathbf{K}_k = \widehat{\mathbf{P}}_k \mathbf{C}_k^T \left(\mathbf{C}_k \widehat{\mathbf{P}}_k \mathbf{C}_k^T + \mathbf{R}_k \right)^{-1} \quad (2.33)$$

Finally, the correction stage yields:

$$\mathbf{x}_k = \widehat{\mathbf{x}}_k + \mathbf{K}_k (\mathbf{z}_k - \mathbf{C}_k \widehat{\mathbf{x}}_k) \quad (2.34)$$

$$\mathbf{P}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \widehat{\mathbf{P}}_k \quad (2.35)$$

Hence, in this section, the KF has been derived by minimizing the variance of the state estimation error and imposing linearity constraints in the formulation of the KF equations. Both the process noise and measurement noise stochastic sequences were assumed to be zero-mean and white but no specific form was imposed for their pdfs. Under the assumption of linear transformations of random variables, the propagation of the mean and covariance are well known and independent of the form of the original pdf itself.

It can readily be seen from (2.34), in the light of its scalar form for simplicity, that as $K_k \rightarrow 0$ (achievable, e.g., if $R \rightarrow +\infty$) the KF privileges the prior state estimation over the measurements whereas the converse applies for $K_k \rightarrow \frac{1}{C_k}$. The KF gain effectively controls the filter's rate of belief in the prediction (theoretical model) or the measurements by weighing the covariance of both. Moreover, the KF gain is independent of the system's state and measurements. It depends only on the system dynamics and the initial guess for the posterior covariance of the state estimation error. It is therefore not surprising

that, in the context of a typical real-time implementation of a time-invariant KF, the gain itself can be pre-calculated either by recursively evolving (2.32), (2.33) and (2.35) or, equivalently, evolving the Riccati equation for the covariance of the prior state estimation error [Kalm 61, Bitt 91] (until some steady-state is achieved) and substituting it in (2.33).

In the following section one will demonstrate how the same set of equations can be derived in the context of a Bayesian formalism.

2.2.2 Bayesian derivation

Let one consider the problem of estimating the state of a system at a discrete time k given all previous noisy measurements up to, and including, k . The conditional pdf of a system being in a state \mathbf{x}_k given all measurements $\mathbf{z}^k = \mathbf{z}_k, \mathbf{z}_{k-1}, \dots, \mathbf{z}_1$ is $p(\mathbf{x}_k|\mathbf{z}^k)$. Using Bayes rule one has (2.36), where $p(\mathbf{x}_k, \mathbf{z}^k)$ is the joint pdf for the system to be in state \mathbf{x}_k and having had the sequence of \mathbf{z}^k measurements.

$$p(\mathbf{x}_k|\mathbf{z}^k) = \frac{p(\mathbf{x}_k, \mathbf{z}^k)}{p(\mathbf{z}^k)} = \frac{p(\mathbf{x}_k, \mathbf{z}_k, \mathbf{z}^{k-1})}{p(\mathbf{z}_k, \mathbf{z}^{k-1})} \quad (2.36)$$

Further application of Bayes rule to the numerator of the previous yields (2.37).

$$p(\mathbf{x}_k, \mathbf{z}_k, \mathbf{z}^{k-1}) = p(\mathbf{z}_k|\mathbf{x}_k, \mathbf{z}^{k-1})p(\mathbf{x}_k, \mathbf{z}^{k-1}) = p(\mathbf{z}_k|\mathbf{x}_k, \mathbf{z}^{k-1})p(\mathbf{x}_k|\mathbf{z}^{k-1})p(\mathbf{z}^{k-1}) \quad (2.37)$$

If, for now, one assumes that, given \mathbf{x}_k , \mathbf{z}_k is conditionally independent of \mathbf{z}^{k-1} , the previous can be simplified to (2.38).

$$p(\mathbf{x}_k, \mathbf{z}_k, \mathbf{z}^{k-1}) = p(\mathbf{z}_k|\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{z}^{k-1})p(\mathbf{z}^{k-1}) \quad (2.38)$$

Substituting the previous in (2.36) and noting that $p(\mathbf{z}_k, \mathbf{z}^{k-1}) = p(\mathbf{z}_k|\mathbf{z}^{k-1})p(\mathbf{z}^{k-1})$ yields (2.39).

$$p(\mathbf{x}_k|\mathbf{z}^k) = \frac{p(\mathbf{z}_k|\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{z}^{k-1})p(\mathbf{z}^{k-1})}{p(\mathbf{z}_k|\mathbf{z}^{k-1})p(\mathbf{z}^{k-1})} \quad (2.39)$$

One can finally write the Bayesian recursive relation for the state conditional pdf.

$$p(\mathbf{x}_k|\mathbf{z}^k) = \frac{p(\mathbf{z}_k|\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{z}^{k-1})}{p(\mathbf{z}_k|\mathbf{z}^{k-1})} \quad (2.40)$$

In the previous, $p(\mathbf{x}_k|\mathbf{z}^{k-1})$ is often called the *prior*, i.e., the conditional pdf that the system's in state \mathbf{x}_k given all measurements up to (but not including) k . It is related with the forecast/prediction stage of the KF. The $p(\mathbf{z}_k|\mathbf{x}_k)$ is often called the *likelihood*, i.e., the conditional pdf that measurements will have the values \mathbf{z}_k given that the system is in state \mathbf{x}_k . The $p(\mathbf{z}_k|\mathbf{z}^{k-1})$ is often called the *evidence* and represents the pdf of a set of measurements \mathbf{z}_k . It is essentially a normalization term. Finally $p(\mathbf{x}_k|\mathbf{z}^k)$ is the *posterior*, i.e., the conditional pdf that the system's in state \mathbf{x}_k given all measurements up to (and including) k .

As in the KF derivation presented in the previous section, one will assume that the system's measurement equation is linear and given by (2.4). Assuming that \mathbf{v}_k is a zero-mean ($E[\mathbf{v}_k] = \mathbf{0}$) stochastic variable drawable from a multivariate normal distribution (white with $E[\mathbf{v}_i\mathbf{v}_j^T] = \delta_{ij}\mathbf{R}_i$) one can write (2.41) and (2.42).

$$E[\mathbf{z}_k|\mathbf{x}_k] = E[\mathbf{C}_k\mathbf{x}_k + \mathbf{v}_k] = \mathbf{C}_k\mathbf{x}_k \quad (2.41)$$

$$\text{Cov} [\mathbf{z}_k | \mathbf{x}_k] = \text{Cov} [\mathbf{v}_k] = \mathbf{R}_k \quad (2.42)$$

Building on the previous assumptions one can materialise $p(\mathbf{z}_k | \mathbf{x}_k)$ yielding (2.43) where $\mathbf{z}_k \in \mathbb{R}^{N_z}$. It is clear from this that the assumption that previously enabled the writing of (2.38) is valid.

$$p(\mathbf{z}_k | \mathbf{x}_k) = \frac{\exp \left[-\frac{1}{2} (\mathbf{z}_k - \mathbf{C}_k \mathbf{x}_k)^T \mathbf{R}_k^{-1} (\mathbf{z}_k - \mathbf{C}_k \mathbf{x}_k) \right]}{(2\pi)^{\frac{N_z}{2}} |\mathbf{R}_k|^{\frac{1}{2}}} \quad (2.43)$$

Similarly as in the KF derivation presented in the previous section, one will assume that the system's process equation is linear and given by (2.44) where \mathbf{x}_k denotes the prior state estimate's random variable and \mathbf{x}_{k-1} denotes the previous posterior state estimate's random variable. Assuming that \mathbf{w}_k is a zero-mean ($E[\mathbf{w}_k] = \mathbf{0}$) stochastic variable drawable from a multivariate normal distribution (white with $E[\mathbf{w}_i \mathbf{w}_j^T] = \delta_{ij} \chi_{ij}$) one can write (2.45) and (2.46).

$$\mathbf{x}_k = \mathbf{A}_{k-1} \mathbf{x}_{k-1} + \mathbf{B}_{k-1} \mathbf{u}_{k-1} + \mathbf{w}_{k-1} \quad (2.44)$$

$$\begin{aligned} E[\mathbf{x}_k | \mathbf{z}^{k-1}] &= E[\mathbf{A}_{k-1} \mathbf{x}_{k-1} + \mathbf{B}_{k-1} \mathbf{u}_{k-1} + \mathbf{w}_{k-1}] \\ &= \mathbf{A}_{k-1} E[\mathbf{x}_{k-1}] + \mathbf{B}_{k-1} \mathbf{u}_{k-1} \\ &= \hat{\mathbf{x}}_k \end{aligned} \quad (2.45)$$

$$\text{Cov} [\mathbf{x}_k | \mathbf{z}^{k-1}] = \hat{\mathbf{P}}_k \quad (2.46)$$

Building on the previous assumptions one can also materialise $p(\mathbf{x}_k | \mathbf{z}^{k-1})$ yielding (2.47) where $\mathbf{x}_k \in \mathbb{R}^{N_x}$.

$$p(\mathbf{x}_k | \mathbf{z}^{k-1}) = \frac{\exp \left[-\frac{1}{2} (\mathbf{x}_k - \hat{\mathbf{x}}_k)^T \hat{\mathbf{P}}_k^{-1} (\mathbf{x}_k - \hat{\mathbf{x}}_k) \right]}{(2\pi)^{\frac{N_x}{2}} |\hat{\mathbf{P}}_k|^{\frac{1}{2}}} \quad (2.47)$$

Substituting (2.43) and (2.47) in (2.40) one can write (2.48) noting that the denominator in (2.40) can be regarded as a normalising constant. Equation (2.48) is therefore an un-normalised posterior pdf.

$$p(\mathbf{x}_k | \mathbf{z}^k) \propto \exp \left[-\frac{1}{2} (\mathbf{z}_k - \mathbf{C}_k \mathbf{x}_k)^T \mathbf{R}_k^{-1} (\mathbf{z}_k - \mathbf{C}_k \mathbf{x}_k) - \frac{1}{2} (\mathbf{x}_k - \hat{\mathbf{x}}_k)^T \hat{\mathbf{P}}^{-1} (\mathbf{x}_k - \hat{\mathbf{x}}_k) \right] \quad (2.48)$$

Let one now find the state \mathbf{x}_k that maximises the posterior pdf given by the previous. This criterion is usually known as the Maximum A Posteriori (MAP). Maximising the previous is equivalent to maximising its natural logarithm. Thus, differentiating and setting it equal to zero yields (2.49).

$$\frac{\partial \ln [p(\mathbf{x}_k | \mathbf{z}^k)]}{\partial \mathbf{x}_k} = \frac{\partial}{\partial \mathbf{x}_k} \left[(\mathbf{z}_k - \mathbf{C}_k \mathbf{x}_k)^T \mathbf{R}_k^{-1} (\mathbf{z}_k - \mathbf{C}_k \mathbf{x}_k) + (\mathbf{x}_k - \hat{\mathbf{x}}_k)^T \hat{\mathbf{P}}^{-1} (\mathbf{x}_k - \hat{\mathbf{x}}_k) \right] = 0 \quad (2.49)$$

Expanding the previous and taking its transpose yields (2.50).

$$2\mathbf{C}_k^T (\mathbf{R}_k^{-1})^T \mathbf{C}_k \mathbf{x}_k + 2 \left(\widehat{\mathbf{P}}_k^{-1} \right)^T \mathbf{x}_k - \mathbf{C}_k^T (\mathbf{R}_k^{-1})^T \mathbf{z}_k - \mathbf{C}_k^T \mathbf{R}_k^{-1} \mathbf{z}_k - \widehat{\mathbf{P}}_k^{-1} \widehat{\mathbf{x}}_k - \left(\widehat{\mathbf{P}}_k^{-1} \right)^T \widehat{\mathbf{x}}_k = \mathbf{0} \quad (2.50)$$

For \mathbf{R}_k and $\widehat{\mathbf{P}}_k$ symmetric one can write (2.51).

$$\mathbf{x}_k = \left(\mathbf{C}_k^T \mathbf{R}_k^{-1} \mathbf{C}_k + \widehat{\mathbf{P}}_k^{-1} \right)^{-1} \times \left(\mathbf{C}_k^T \mathbf{R}_k^{-1} \mathbf{z}_k + \widehat{\mathbf{P}}_k^{-1} \widehat{\mathbf{x}}_k \right) \quad (2.51)$$

Assuming that \mathbf{R}_k and $\widehat{\mathbf{P}}_k$ are invertible and using the matrix inversion lemma [Hend 81] to calculate the matrix inversion in the previous equation, one can write (2.52).

$$\mathbf{x}_k = \left[\widehat{\mathbf{P}}_k - \widehat{\mathbf{P}}_k \mathbf{C}_k^T \left(\mathbf{R}_k + \mathbf{C}_k \widehat{\mathbf{P}}_k \mathbf{C}_k^T \right)^{-1} \mathbf{C}_k \widehat{\mathbf{P}}_k \right] \times \left[\widehat{\mathbf{P}}_k^{-1} \widehat{\mathbf{x}}_k + \mathbf{C}_k^T \mathbf{R}_k^{-1} \mathbf{z}_k \right] \quad (2.52)$$

Expanding the previous and defining $\mathbf{K}_k = \widehat{\mathbf{P}}_k \mathbf{C}_k^T \left(\mathbf{C}_k \widehat{\mathbf{P}}_k \mathbf{C}_k^T + \mathbf{R}_k \right)^{-1}$ gives finally the well known KF state update equation (2.53).

$$\mathbf{x}_k = \widehat{\mathbf{x}}_k + \mathbf{K}_k (\mathbf{z}_k - \mathbf{C}_k \widehat{\mathbf{x}}_k) \quad (2.53)$$

Hence, \mathbf{x}_k is the estimate that maximizes the state posterior pdf, given the history of measurements up to k . At this point one has no knowledge of the form of this pdf or even its covariance. In order to find these out let one start by writing (2.53) as (2.54). \mathbf{x}_k is therefore explicitly written as a linear combination of $\widehat{\mathbf{x}}_k$ and \mathbf{z}_k , both normally distributed stochastic variables.

$$\mathbf{x}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \widehat{\mathbf{x}}_k + \mathbf{K}_k \mathbf{z}_k \quad (2.54)$$

Definition Let $\boldsymbol{\alpha} \sim \mathcal{N}(\boldsymbol{\mu}_\alpha, \boldsymbol{\Sigma}_\alpha)$ for which the moment generating function is known to be $\mathbf{M}_\alpha(\mathbf{t}) = E[\exp(\mathbf{t}^T \boldsymbol{\alpha})] = \exp[\boldsymbol{\mu}_\alpha^T \mathbf{t} + \frac{1}{2} \mathbf{t}^T \boldsymbol{\Sigma}_\alpha \mathbf{t}]$. In addition, let the random variable $\boldsymbol{\beta}$ be defined as $\boldsymbol{\beta} = \boldsymbol{\Lambda} \boldsymbol{\alpha}$.

Lemma 2.2.1 $\boldsymbol{\beta}$ is distributed such that $\boldsymbol{\beta} \sim \mathcal{N}(\boldsymbol{\Lambda} \boldsymbol{\mu}_\alpha, \boldsymbol{\Lambda} \boldsymbol{\Sigma}_\alpha \boldsymbol{\Lambda}^T)$.

Proof

$$\begin{aligned} \mathbf{M}_\beta(\mathbf{t}) &= E[\exp(\mathbf{t}^T \boldsymbol{\beta})] = E[\exp(\mathbf{t}^T \boldsymbol{\Lambda} \boldsymbol{\alpha})] \\ &= \exp \left[\boldsymbol{\mu}_\alpha^T (\mathbf{t}^T \boldsymbol{\Lambda})^T + \frac{1}{2} \mathbf{t}^T \boldsymbol{\Lambda} \boldsymbol{\Sigma}_\alpha (\mathbf{t}^T \boldsymbol{\Lambda})^T \right] \\ &= \exp \left[(\boldsymbol{\Lambda} \boldsymbol{\mu}_\alpha)^T \mathbf{t} + \frac{1}{2} \mathbf{t}^T \boldsymbol{\Lambda} \boldsymbol{\Sigma}_\alpha \boldsymbol{\Lambda}^T \mathbf{t} \right] \end{aligned} \quad (2.55)$$

Hence $\boldsymbol{\beta} \sim \mathcal{N}(\boldsymbol{\Lambda} \boldsymbol{\mu}_\alpha, \boldsymbol{\Lambda} \boldsymbol{\Sigma}_\alpha \boldsymbol{\Lambda}^T)$.

Hence, any linear transformation of a normally distributed random variable is itself a normally distributed random variable.

Definition Let $\boldsymbol{\gamma} = \boldsymbol{\Upsilon} \boldsymbol{\zeta} + \boldsymbol{\Xi} \boldsymbol{\eta}$, where $\boldsymbol{\zeta} \sim \mathcal{N}(\boldsymbol{\mu}_\zeta, \boldsymbol{\Sigma}_\zeta)$ and $\boldsymbol{\eta} \sim \mathcal{N}(\boldsymbol{\mu}_\eta, \boldsymbol{\Sigma}_\eta)$ are independent random variables.

Lemma 2.2.2 $\boldsymbol{\gamma}$ is distributed such that $\boldsymbol{\gamma} \sim \mathcal{N}(\boldsymbol{\Upsilon} \boldsymbol{\mu}_\zeta + \boldsymbol{\Xi} \boldsymbol{\mu}_\eta, \boldsymbol{\Upsilon} \boldsymbol{\Sigma}_\zeta \boldsymbol{\Upsilon}^T + \boldsymbol{\Xi} \boldsymbol{\Sigma}_\eta \boldsymbol{\Xi}^T)$.

Proof Using lemma 2.2.1, one has that $(\Upsilon\zeta) \sim \mathcal{N}(\Upsilon\mu_\zeta, \Upsilon\Sigma_\zeta\Upsilon^T)$ and $(\Xi\eta) \sim \mathcal{N}(\Xi\mu_\eta, \Xi\Sigma_\eta\Xi^T)$. Because ζ and η are independent, hence $(\Upsilon\zeta)$ and $(\Xi\eta)$ are also independent, a basic property of moment generating functions yields $\mathbf{M}_\gamma(\mathbf{t}) = \mathbf{M}_{(\Upsilon\zeta)}(\mathbf{t}) \times \mathbf{M}_{(\Xi\eta)}(\mathbf{t})$. So,

$$\begin{aligned} \mathbf{M}_\gamma(\mathbf{t}) &= \exp\left[(\Upsilon\mu_\zeta)^T \mathbf{t} + \frac{1}{2}\mathbf{t}^T \Upsilon\Sigma_\zeta\Upsilon^T \mathbf{t}\right] \times \exp\left[(\Xi\mu_\eta)^T \mathbf{t} + \frac{1}{2}\mathbf{t}^T \Xi\Sigma_\eta\Xi^T \mathbf{t}\right] \\ &= \exp\left[(\Upsilon\mu_\zeta + \Xi\mu_\eta)^T \mathbf{t} + \frac{1}{2}\mathbf{t}^T (\Upsilon\Sigma_\zeta\Upsilon^T + \Xi\Sigma_\eta\Xi^T) \mathbf{t}\right] \end{aligned} \quad (2.56)$$

hence $\gamma \sim \mathcal{N}(\Upsilon\mu_\zeta + \Xi\mu_\eta, \Upsilon\Sigma_\zeta\Upsilon^T + \Xi\Sigma_\eta\Xi^T)$.

So, returning to (2.54), one has shown that $\mathbf{x}_k \sim \mathcal{N}\left(\widehat{\mathbf{x}}_k + \mathbf{K}_k(\mathbf{z}_k - \mathbf{C}_k\widehat{\mathbf{x}}_k), (\mathbf{I} - \mathbf{K}_k\mathbf{C}_k)\widehat{\mathbf{P}}_k(\mathbf{I} - \mathbf{K}_k\mathbf{C}_k)^T + \mathbf{K}_k\Sigma\mathbf{K}_k^T\right)$ where the random variables (RV) $\widehat{\mathbf{x}}_k^{RV} \sim \mathcal{N}\left(\widehat{\mathbf{x}}_k, \widehat{\mathbf{P}}_k\right)$ and $\mathbf{z}_k^{RV} \sim \mathcal{N}(\mathbf{C}_k\mathbf{x}_k^{true}, \Sigma)$. Since the covariance of a random variable is invariant to translations, one can define $\mathbf{P}_k = (\mathbf{I} - \mathbf{K}_k\mathbf{C}_k)\widehat{\mathbf{P}}_k(\mathbf{I} - \mathbf{K}_k\mathbf{C}_k)^T + \mathbf{K}_k\Sigma\mathbf{K}_k^T$, the covariance of the posterior state estimation error, and substitute \mathbf{K}_k , from (2.33), to give (2.35).

Finally, one needs to express $\widehat{\mathbf{P}}_k$ as a function of \mathbf{P}_{k-1} in order to obtain the full set of KF equations. Because the present prior state estimate is a linear transformation of the previous posterior state estimate, equation (2.44), if the latter's covariance is \mathbf{P}_{k-1} then, from lemma 2.2.1, the former's covariance is $\widehat{\mathbf{P}}_k = \mathbf{A}_{k-1}\mathbf{P}_{k-1}\mathbf{A}_{k-1}^T + \mathbf{Q}_{k-1}$. Once again, because the covariance is invariant to translations, $\widehat{\mathbf{P}}_k$ is also the covariance of the prior state estimation error.

At this point it is worth noting that one has derived the same set of KF equations in two different contexts. Clearly the assumptions required to derive the KF equations in this section are a more stringent than the ones required in section 2.2.1, however, if satisfied, it is guaranteed that the state estimate provided by the KF is the mode of the conditional posterior pdf. If, for instance, the normality assumptions are not valid but the white and zero-mean assumptions are, then the estimate may no longer be the mode of the conditional posterior pdf but it is still the one with the MMSE.

2.2.3 Scalar time-invariant Kalman filter gain

In this section one will demonstrate how to calculate the steady-state gain in the case of a scalar, linear and time-invariant Kalman filter implementation. Let one start by writing the celebrated Kalman filter equations (2.31) to (2.35) in the particular case of a scalar linear time-invariant implementation thus yielding equations (2.57) to (2.61).

$$\widehat{x}_k = ax_{k-1} + bu_{k-1} \quad (2.57)$$

$$\widehat{P}_k = a^2 P_{k-1} + Q \quad (2.58)$$

$$K_k = c\widehat{P}_k \left(c^2\widehat{P}_k + R\right)^{-1} \quad (2.59)$$

$$x_k = \widehat{x}_k + K_k(z_k - c\widehat{x}_k) \quad (2.60)$$

$$P_k = (1 - cK_k)\hat{P}_k \quad (2.61)$$

Let one substitute (2.61) in (2.58) to write (2.62).

$$\hat{P}_k = a^2(1 - cK_{k-1})\hat{P}_{k-1} + Q \quad (2.62)$$

In steady-state conditions one has that the steady-state covariance of the prior state estimation error (\hat{P}_{SS}) satisfies $\hat{P}_{SS} \approx \hat{P}_k \approx \hat{P}_{k-1}$ thus enabling one to write (2.63).

$$\hat{P}_{SS} = \hat{P}_k = Q(1 - a^2 + ca^2K_{k-1})^{-1} \quad (2.63)$$

Substituting the previous in (2.59) yields (2.64).

$$K_k = cQ [c^2Q + (1 - a^2 + ca^2K_{k-1})R]^{-1} \quad (2.64)$$

Again, in steady-state conditions, the steady-state Kalman gain (K_{SS}) satisfies $K_{SS} \approx K_k \approx K_{k-1}$ thus enabling one to write (2.65).

$$ca^2RK_{SS}^2 + [R(1 - a^2) + c^2Q]K_{SS} - cQ = 0 \quad (2.65)$$

This is a quadratic equation which admits two solutions. It is straightforward to realize that, as far as the Kalman filter is concerned, there is only one meaningful solution which must lie in the interval $[0 \ c^{-1}]$ if $c > 0$ or $[c^{-1} \ 0]$ if $c < 0$.

Figure 2.1 shows a contour plot of the Kalman filter gain as a function of the process variance Q and measurement variance R for the case where $a = c = 1$ independently of b .

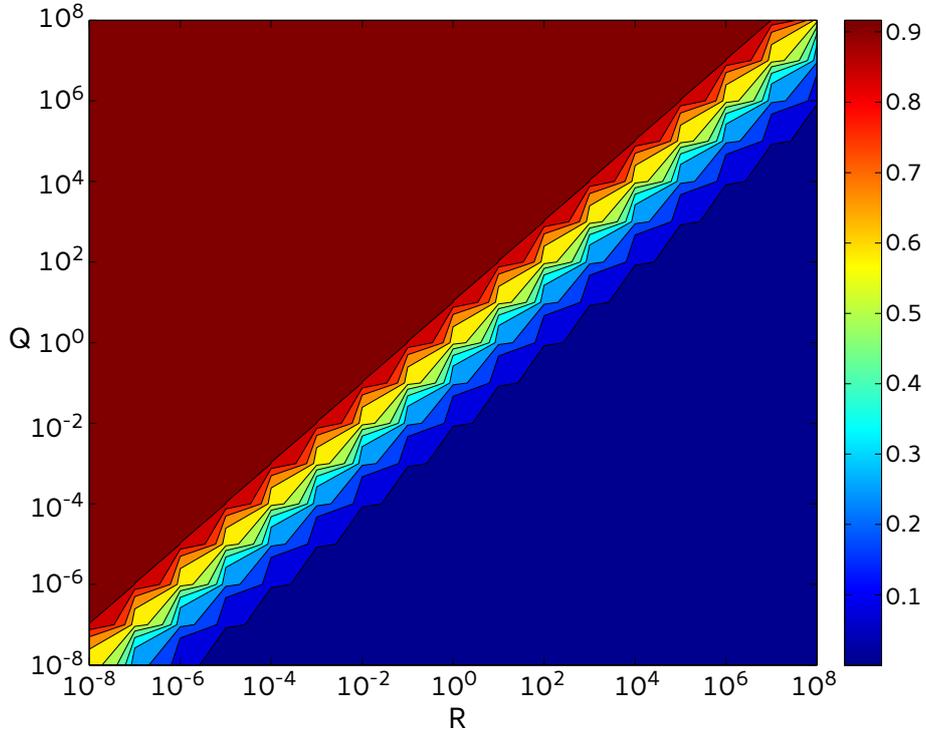


Figure 2.1: Steady-state time-invariant scalar Kalman filter gain contour as a function of R and Q

It is also worthwhile noting that, in the particular case of $a = c = 1$, and for $R \gg Q$, the steady-state Kalman filter gain K_{SS} converges asymptotically to $\sqrt{Q/R}$.

2.3 Non-linear Kalman filters

In this section, several non-linear implementations of the KF are described including a comprehensive derivation of the EKF, the basis of them all.

2.3.1 Extended Kalman filter

The EKF is the most popular and the eldest of all KF evolutions addressing state estimation of systems exhibiting non-linear dynamics. Arguably, its main advantage over other KF-based non-linear approaches is its simplicity. Considering a system for which one has a set of generic time-dependent non-linear state equations of the form shown in equation (2.66) and where one assumes, for simplicity and without loss of generality, no inputs.

$$\Psi_k = \mathbf{f}_{k-1}(\Psi_{k-1}) \quad (2.66)$$

If one also assumes that Ψ_{k-1} is in the neighbourhood of some state, $\tilde{\Psi}_{k-1}$, the previous can be linearized to obtain equation (2.67), where $\left[\frac{\partial \mathbf{f}_{k-1}}{\partial \Psi}\right]_{\tilde{\Psi}_{k-1}}$ is the Jacobian of \mathbf{f}_{k-1} evaluated at $\tilde{\Psi}_{k-1}$.

$$\Psi_k \approx \mathbf{f}_{k-1}(\tilde{\Psi}_{k-1}) + \left[\frac{\partial \mathbf{f}_{k-1}}{\partial \Psi}\right]_{\tilde{\Psi}_{k-1}} \times (\Psi_{k-1} - \tilde{\Psi}_{k-1}) \quad (2.67)$$

Equations (2.68) and (2.69) allow the rearrangement of equation (2.67) yielding the linear approximation shown in equation (2.70).

$$\mathbf{F}_{k-1} = \left[\frac{\partial \mathbf{f}_{k-1}}{\partial \Psi}\right]_{\tilde{\Psi}_{k-1}} \quad (2.68)$$

$$\mathbf{U}_{k-1} = \mathbf{f}_{k-1}(\tilde{\Psi}_{k-1}) - \left[\frac{\partial \mathbf{f}_{k-1}}{\partial \Psi}\right]_{\tilde{\Psi}_{k-1}} \times \tilde{\Psi}_{k-1} \quad (2.69)$$

$$\Psi_k \approx \mathbf{F}_{k-1} \Psi_{k-1} + \mathbf{U}_{k-1} \quad (2.70)$$

Having found a linearized expression for the state equation allows us to write the predictor stage covariance propagation for the EKF in analogy with equation (2.32) by letting $\mathbf{A}_k = \mathbf{F}_{k-1}$. Whereas in the case of the linear KF this relation is exact, in the case of the EKF it is approximate.

$$\hat{\mathbf{P}}_k = \mathbf{F}_{k-1} \mathbf{P}_{k-1} \mathbf{F}_{k-1}^T + \mathbf{Q}_{k-1} \quad (2.71)$$

Conveniently, the Jacobian (\mathbf{F}_{k-1}) is evaluated at $\tilde{\Psi}_{k-1} = \mathbf{x}_{k-1}$, i.e., at the last available posterior state estimate. The state equation itself, equation (2.72), remains non-linear and is used as so.

$$\hat{\mathbf{x}}_k = \mathbf{f}_{k-1}(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}, \mathbf{w}_{k-1}) \quad (2.72)$$

Similarly if the system has a known, potentially non-linear, measurement relation given by equation (2.73), linearizing it about some state, $\tilde{\Psi}_k$ yields equation (2.74).

$$\mathbf{z}_k = \mathbf{h}_k(\Psi_k) \quad (2.73)$$

$$\mathbf{z}_k \approx \mathbf{h}(\tilde{\Psi}_k) + (\Psi_k - \left[\frac{\partial \mathbf{h}_k}{\partial \Psi} \right]_{\tilde{\Psi}_k} \times \tilde{\Psi}_k) \quad (2.74)$$

Rearranging the previous using equations (2.75) and (2.76) yields the linear approximation of the measurement relation shown in equation (2.77).

$$\mathbf{Z}_k = \mathbf{z}_k - \mathbf{h}(\tilde{\Psi}_k) + \left[\frac{\partial \mathbf{h}_k}{\partial \Psi} \right]_{\tilde{\Psi}_k} \times \tilde{\Psi}_k \quad (2.75)$$

$$\mathbf{H}_k = \left[\frac{\partial \mathbf{h}_k}{\partial \Psi} \right]_{\tilde{\Psi}_k} \quad (2.76)$$

$$\mathbf{Z}_k \approx \mathbf{H}_k \Psi_k \quad (2.77)$$

Having this approximate linear relation, and in analogy with the linear KF, equations (2.33) and (2.35) can also be used in the EKF by letting $\mathbf{C}_k = \mathbf{H}_k$. Of course, whereas in the case of the linear KF these are exact, in the case of the EKF they are approximate.

Conveniently, the Jacobian (\mathbf{H}_k) is evaluated at $\tilde{\Psi}_k = \hat{\mathbf{x}}_k$, i.e., at the last available prior state estimate. The measurement relation itself remains non-linear and is used as so for the posterior state estimation.

$$\mathbf{x}_k = \hat{\mathbf{x}}_k + \mathbf{K}_k (\mathbf{z}_k - \mathbf{h}_k(\hat{\mathbf{x}}_k)) \quad (2.78)$$

The linearization of the system dynamics, allowing the usage of the linear KF equations for the propagation of both the prior and posterior covariances, is effectively equivalent to approximating the state's prior and posterior pdfs to normal distributions. Furthermore, and rather conveniently, the linearizations involved in the EKF are performed around state estimates that represent the best available information about the system at each given time. It is straightforward to verify that in case the system dynamics is linear the EKF reverts to the linear KF. Unlike the linear KF though, the EKF is not optimal and may quickly diverge due to linearization errors when applied to highly non-linear system dynamics. This is a consequence of the fact that, again unlike the linear KF, the gain ends up depending on the system measurements themselves via the linearization process. This is clear since the gain depends on the Jacobian of the measurement transformation evaluated at the prior state estimate which depends on the posterior state estimate which, in turn, depends on the measurements. Over the years, several approaches have been proposed to improve the performance and robustness of the EKF. Amongst the most popular ones are the iterated Extended Kalman Filter (iEKF) [Lefe 01] and the robust Extended Kalman Filter (rEKF) [Eini 99, Xion 08].

2.3.2 Iterated extended Kalman filter

The iEKF [Simo 06] attempts to iteratively refine the state estimation around which the Taylor expansion of the measurement equation is done. The idea is to use the ordinary posterior state estimate of the EKF, re-linearize the measurement equation around it and then reevaluate equations (2.33) to (2.35), producing a new posterior state estimate. The operation can be repeated as many times as required until some predefined criteria is met. Obviously, if a system model employs a linear measurement equation, the iEKF is of no added value and in fact reverts to the EKF.

2.3.3 Robust extended Kalman filter

The rEKF, derived using the linear H_∞ technique, is designed to minimize the relative worst case effects of disturbances and linearization errors. While the standard KF and EKF minimize the mean square error of the posterior state estimate, the rEKF tries to limit the normalized energy of the estimation error. Instead of neglecting higher order terms, as in the standard EKF approach, the rEKF acknowledges their existence, treats them as disturbances and minimizes their effect. This is expressed in equation (2.79) where $\|\cdot\|$ denotes the 2-norm, $\tilde{\mathbf{x}}_k$ is the error of the posterior state estimate, \mathbf{w}_k and \mathbf{v}_k are respectively the model and measurement noise, $\mathbf{\Delta}_k$ is the linearization error (expansion terms of order > 1 ignored in equation (2.67)) and γ is a design parameter. It can be shown [Xion 08] that, if γ satisfies some (mostly bounding) conditions, the estimation error of the rEKF satisfies equation (2.79).

$$\frac{\|\tilde{\mathbf{x}}_k\|^2}{\|\mathbf{w}_k\|^2 + \|\mathbf{v}_k\|^2 + \|\mathbf{\Delta}_k\|^2} \leq \gamma^2 \quad (2.79)$$

Practically, the lower the γ the more confined the normalized estimation error and consequently the more stable/robust the filter. The full set of equations for the rEKF with a linear measurement relation is given by equations (2.34), (2.71), (2.72) and (2.80) to (2.82).

$$\hat{\mathbf{\Sigma}}_k = \left(\hat{\mathbf{P}}_k^{-1} - \gamma^{-2} \mathbf{I} \right)^{-1} \quad (2.80)$$

$$\mathbf{K}_k = \hat{\mathbf{\Sigma}}_k \mathbf{C}_k^T \left(\mathbf{C}_k \hat{\mathbf{\Sigma}}_k \mathbf{C}_k^T + \mathbf{R}_k \right)^{-1} \quad (2.81)$$

$$\mathbf{P}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_k) \hat{\mathbf{\Sigma}}_k \quad (2.82)$$

The predictor stage of the rEKF is basically given by the EKF equations but now a different prior covariance is calculated based on the EKF's standard one. The corrector stage of the rEKF is given by the linear KF equations with this new prior covariance replacing the standard one. Interestingly, and as expected, $\lim_{\gamma \rightarrow \infty} \hat{\mathbf{\Sigma}}_k = \hat{\mathbf{P}}_k$, i.e., the rEKF reverts to the EKF for large values of γ leaving the normalized estimation error unbounded. This means that the γ parameter is effectively a trade-off knob between optimality and stability.

The rEKF is, as will be demonstrated in this work, a very effective refinement of the EKF, exhibiting the same level of complexity and very little additional computational cost.

2.3.4 Unscented Kalman filter

The EKF works by linearising the transformation undergone by a normally distributed random variable. This is equivalent to approximating the resultant pdf also to a normal distribution. Therefore, assuming Gaussian distributions up and downstream, the propagation of the mean and covariance is well established. However, in the presence of highly non-linear dynamics, cumulative errors arising from this linearisation process often have drastic consequences on the EKF's stability.

The Unscented Kalman Filter (UKF) [Juli 97] is a widely acknowledged improvement over EKF-based methods. Instead of assuming a shape for the pdf and/or approximating the state/measurement transformation, the UKF uses the Unscented Transform (UT)

to propagate the mean and covariance without any knowledge whatsoever about the underlying pdf.

Brief introduction to the unscented transform

The underlying idea behind the UT is to take a set of random variable values (sigma points) representative of a pdf and propagate those through the non-linear transformation. This set is chosen deterministically based on the mean and covariance known at the time. After the propagation process, the sample mean and sample covariance values are updated and a new set of sigma points can be calculated. Unlike the linearization approach, the UT provides accurate estimates up to third order in the usual Taylor expansion [Simo 06].

The UKF implementation

The first step is to use the relations of equation (2.83) to calculate the set of $2N + 1$ (where N is the dimension of the state vector) sigma points [Juli 97] representative of the posterior state estimate pdf for the previous time sample. In these, $\tilde{\mathbf{x}}_{k-1}^i$ is a sigma point, \mathbf{x}_{k-1} is the posterior state estimate for the previous time sample, P_{k-1} is the posterior covariance estimate for the previous time sample, λ is a scalar quantity which will be discussed a bit later and $\left(\sqrt{(N + \lambda)P_{k-1}}\right)_i$ is the i^{th} row of the Cholesky decomposition [Horn 90] of $(N + \lambda)P_{k-1}$.

$$\begin{aligned}\tilde{\mathbf{x}}_{k-1}^0 &= \mathbf{x}_{k-1} \\ \tilde{\mathbf{x}}_{k-1}^i &= \mathbf{x}_{k-1} + \left(\sqrt{(N + \lambda)P_{k-1}}\right)_i^T, \quad i = 1, \dots, N \\ \tilde{\mathbf{x}}_{k-1}^i &= \mathbf{x}_{k-1} - \left(\sqrt{(N + \lambda)P_{k-1}}\right)_{i-N}^T, \quad i = N + 1, \dots, 2N\end{aligned}\quad (2.83)$$

These sigma points are then propagated through the state equation (2.84), where \mathbf{u}_k represents the system inputs, to produce a new set, $\hat{\mathbf{x}}_k^i$, representative of the prior state estimate pdf for the present time sample.

$$\hat{\mathbf{x}}_k^i = \mathbf{f}_{k-1}(\tilde{\mathbf{x}}_{k-1}^i, \mathbf{u}_k) \quad (2.84)$$

With these, modified versions of the usual sample mean, equation (2.85), and sample covariance, equation (2.86), formulas are used to calculate the new pdf's first, $\hat{\mathbf{x}}_k$, and second, $\hat{\mathbf{P}}_k^x$, moments. The \mathbf{Q} matrix represents the state model's covariance and the weights, W^i , are given by the relations shown in equation (2.87) with the scaling factor $\lambda = \alpha^2(N + \kappa) - N$, where α determines the spread of the sigma points around the mean, κ is yet another scaling parameter and β is used to incorporate prior pdf knowledge [Wan 00], $\beta = 2$ is optimal for Gaussian distributions.

$$\hat{\mathbf{x}}_k = \sum_{i=0}^{2N} W_{mean}^i \hat{\mathbf{x}}_k^i \quad (2.85)$$

$$\hat{\mathbf{P}}_k^x = \sum_{i=0}^{2N} W_{covar}^i (\hat{\mathbf{x}}_k^i - \hat{\mathbf{x}}_k)(\hat{\mathbf{x}}_k^i - \hat{\mathbf{x}}_k)^T + \mathbf{Q} \quad (2.86)$$

$$\begin{aligned}W_{mean}^0 &= \frac{\lambda}{N + \lambda} \\ W_{covar}^0 &= \frac{\lambda}{N + \lambda} + (1 - \alpha^2 + \beta) \\ W_{mean}^i &= W_{covar}^i = \frac{1}{2(N + \lambda)}, \quad i = 1, \dots, 2N\end{aligned}\quad (2.87)$$

Having now an approximate characterization of the predicted state's pdf, one moves on to get an approximate characterization of the predicted measurement's pdf using the exact same procedure. The sigma points are first calculated using the relations of equation (2.88) and propagated through the measurement equation (2.89).

$$\begin{aligned}\tilde{\mathbf{y}}_k^0 &= \hat{\mathbf{x}}_k \\ \tilde{\mathbf{y}}_k^i &= \hat{\mathbf{x}}_k + \left(\sqrt{(N + \lambda) \hat{\mathbf{P}}_k^x} \right)_i^T, \quad i = 1, \dots, N \\ \tilde{\mathbf{y}}_k^i &= \hat{\mathbf{x}}_k - \left(\sqrt{(N + \lambda) \hat{\mathbf{P}}_k^x} \right)_{i-N}^T, \quad i = N + 1, \dots, 2N\end{aligned}\tag{2.88}$$

$$\hat{\mathbf{y}}_k^i = \mathbf{h}_{k-1}(\tilde{\mathbf{y}}_k^i)\tag{2.89}$$

The new pdf is characterized by using modified versions of the usual sample mean, equation (2.90), and sample covariance, equation (2.91), to calculate its first two moments: $\hat{\mathbf{y}}_k$ and \mathbf{P}_k^y , respectively. The \mathbf{R} matrix represents the measurement's covariance.

$$\hat{\mathbf{y}}_k = \sum_{i=0}^{2N} W_{mean}^i \hat{\mathbf{y}}_k^i\tag{2.90}$$

$$\mathbf{P}_k^y = \sum_{i=0}^{2N} W_{covar}^i (\hat{\mathbf{y}}_k^i - \hat{\mathbf{y}}_k)(\hat{\mathbf{y}}_k^i - \hat{\mathbf{y}}_k)^T + \mathbf{R}\tag{2.91}$$

The KF gain is calculated using equation (2.92) where \mathbf{P}_k^{xy} , see equation (2.93), is the state and measurement cross covariance matrix. It can be shown [Simo 06] that equation (2.92) is equivalent to equation (2.33).

$$\mathbf{K}_k = \mathbf{P}_k^{xy} (\mathbf{P}_k^y)^{-1}\tag{2.92}$$

$$\mathbf{P}_k^{xy} = \sum_{i=0}^{2N} W_{covar}^i (\hat{\mathbf{x}}_k^i - \hat{\mathbf{x}}_k)(\hat{\mathbf{y}}_k^i - \hat{\mathbf{y}}_k)^T\tag{2.93}$$

Finally, the posterior state and covariance estimates are given, respectively, by equations (2.94) and (2.95) [Simo 06].

$$\mathbf{x}_k = \hat{\mathbf{x}}_k + \mathbf{K}_k (\mathbf{y}_k - \hat{\mathbf{y}}_k)\tag{2.94}$$

$$\mathbf{P}_k = \hat{\mathbf{P}}_k - \mathbf{K}_k \mathbf{P}_k^y \mathbf{K}_k^T\tag{2.95}$$

Having introduced the UKF, and having covered the basics of the types of KFs used in this work, one now moves on to discuss the system dynamics and the formulation of the problem to be solved.

Chapter 3

Linear harmonic estimation

3.1 Introduction

The requirements for accurate analysis of instantaneous quantities associated with signal components, earlier in the analog domain and nowadays also in the digital domain, has challenged engineers and scientists of many scientific areas. Requirements vary depending on the particular application but might involve, for example, harmonic amplitude growth detection [Saiz 97], frequency tracking [Chic 94] or distortion estimation [Beid 91]. In many systems, automated control schemes are devised for acting accordingly based on real-time estimations. It is important therefore, in such systems, to minimize estimation latencies so that the control action can be opportune and more effective. Furthermore, with the advent of modern control theory, taking advantage and defining highly complex multiple-input and multiple-output systems making use of the state-space formalism, the capability of simultaneously (and accurately) estimating various quantities in real-time is of the essence.

In tokamak fusion research, KF [Kalm 60] based methods have already proven themselves efficient in numerous applications. They have, for instance, been used in the estimation of the electron density profile from reflectometry diagnostic data [Nune 99], as part of control schemes for stabilizing resistive wall modes [In 06a] and for vessel current estimations [Ou 07] in DIII-D, in the real-time estimation of the magnetic field pitch angle from motional Stark effect diagnostic data [Coel 09a, Coel 09b] in JET, for real-time mode number and spectral analysis [Coel 08] in ISTTOK and for synchronous detection and resonance identification with the toroidal Alfvén eigenmodes antenna diagnostic [Alve 10] in JET.

KF based methods optimize the state estimation of a system based both on its theoretical model and on observable quantities with known relations with the system state. It is very much in this capability of incorporating information from both theory and measurements, and symbiotically providing the optimal output, that resides its main strength. These methods provide, in a balanced way, the best of both worlds: the deterministic prediction of the system's mathematical model and nature's insubordinate non-ideal behaviour.

One of the greatest merits of the KF relies on its capability in tackling any sort of estimation/filtering problem for which an exact linear state process evolution equation exists. It was not long before the community realised its potential in the field of spectral analysis [Bitm 86]. Indeed, it can be shown that a quasi-periodic signal, sampled at a regular sampling frequency, can be described by a constant increment in the in-

stantaneous phase, related to the instantaneous frequency. As such, devising a process matrix for a KF implementation is fairly straightforward. One may refer to such implementations, and all their derivatives, as the KFHE, the first of which was proposed in [Bitm 86]. In the context of fusion applications, and for the present dissertation, it was used in [Coel 09a, Coel 09b] in a lock-in amplifier [Mich 41] implementation as a real-time amplitude demodulation technique for the processing of data from the motional Stark effect diagnostic [Hawk 99, Hawk 06, Alve 04] at the JET. Later it was used for the synchronous detection of the plasma response (resonant behaviour observable in terms of amplitude and phase) to magnetic perturbations in the Alfvén frequency range in the same tokamak [Alve 10]. It had already been demonstrated to be a real-time capable algorithm in the estimation of the poloidal mode number of MagnetoHydroDynamic (MHD) activity in the ISTTOK tokamak [Coel 08] and finally it was also used as a part of a larger KF-based scheme for the real-time tracking and identification of MHD activity in the JET tokamak [Alve 13b].

This chapter presents a brief introduction to the theory behind the KFHE followed by the exposure of three different applications in the context of nuclear fusion research.

Developments presented herein are based on [Coel 09a, Coel 09b, Coel 08, Alve 10] for which the author of the present dissertation provided a substantial contribution.

3.2 Theory

Assuming a generic signal S_k , see equation (3.1) where k denotes the time index, to be the sum of N non-stationary signal components $S_{i,k}$ exhibiting the properties of an intrinsic mode function [Huan 98], i.e., $S_{i,k} = A_{i,k} \sin(\phi_{i,k})$ (where $A_{i,k}$ is the instantaneous amplitude and $\phi_{i,k}$ is the instantaneous phase) plus a noise term n_k , then \mathbf{f}_k describes the linear harmonic estimator model/transformation for a single signal component, see (3.2). Rewriting the previous in state space notation yields equation (3.3), where $\hat{\mathbf{x}}$ is the prior state estimate and \mathbf{x} is the posterior state estimate, illustrating that the model is a rotation in the I - Q plane with phase increment $\Delta\phi_k = 2\pi f_k/F_s$, where F_s is the signal's sampling frequency. Typically, and in all implementations presented in the present chapter, $\Delta\phi_k$ may be time dependent but it is assumed to be known at all times.

$$S_k = \sum_{i=1}^N S_{i,k} + n_k \quad (3.1)$$

$$\mathbf{f}_k : \mathbb{R}^2 \rightarrow \mathbb{R}^2 : (x_{1,k}, x_{2,k}) \rightarrow \begin{pmatrix} x_{1,k-1} \cos(\Delta\phi_{k-1}) - x_{2,k-1} \sin(\Delta\phi_{k-1}) \\ x_{1,k-1} \sin(\Delta\phi_{k-1}) + x_{2,k-1} \cos(\Delta\phi_{k-1}) \end{pmatrix} \quad (3.2)$$

$$\begin{bmatrix} \hat{x}_{1,k} \\ \hat{x}_{2,k} \end{bmatrix} = \begin{bmatrix} \cos(\Delta\phi_{k-1}) & -\sin(\Delta\phi_{k-1}) \\ \sin(\Delta\phi_{k-1}) & \cos(\Delta\phi_{k-1}) \end{bmatrix} \begin{bmatrix} x_{1,k-1} \\ x_{2,k-1} \end{bmatrix} \quad (3.3)$$

Again, in the case of a single tone, $x_{1,k}$ represents the signal estimate and $x_{2,k}$ can be viewed as the real-time Hilbert transform of $x_{1,k}$ (i.e. the quadrature of $x_{1,k}$). In some senses $x_{1,k}$ can be interpreted as the result of a non-stationary band-pass filter around the tracking signal component. Furthermore, if the only observable quantity is a direct measurement of the signal S_k , then the measurement equation (i.e. the predicted measurement's equation) is linear and, for a single tone, given by equation (3.4).

$$\hat{z}_k = \mathbf{C}_k \hat{\mathbf{x}}_k = [1 \ 0] [\hat{x}_{1,k} \ \hat{x}_{2,k}]^T \quad (3.4)$$

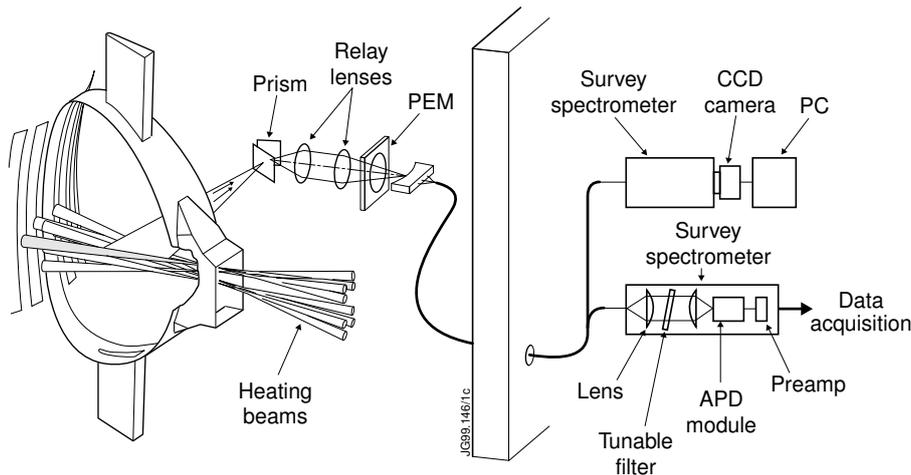


Figure 3.1: Overview of the MSE diagnostic apparatus at JET. *Courtesy of EFDA-JET.*

benefits over other more standard, well established and widely used Fourier based techniques such as the Single Phase Lock-in (SPL) [Shi 06]. These advantages are particularly evident in cases where a limited amount of data samples is available for block processing in typical multi-rate implementations and in the case of frequency jitter in the modulation. Furthermore, the method is also shown to provide significant improvements in terms of the estimates' signal-to-noise ratio by providing fairly straightforward means of attenuating the contaminating effect of the power grid's signal component as well as the typical bursty¹ nature of the plasma's ELMs.

Introduction

KF-based approaches for inferring the harmonic/frequency content of signals in industrial and experimental environments have been extensively explored [Rout 02, Ma 96, Agha 05, Dash 00, Bitt 00, Caci 96, Bitm 86, Grub 94, La S 96]. Ultimately these signals are noisy combinations of non-stationary components thus very hard to individually characterise and discriminate. The MSE diagnostic, widely used in tokamak-based nuclear fusion research with neutral beam injection capabilities, is an important diagnostic as it provides localized measurements of the poloidal field hence assisting in the characterization of the safety factor and current profiles [Hirs 94, Gian 04]. At JET, an overview of the MSE diagnostic apparatus is shown in figure 3.1.

The Stark effect arises from the local electric field generated by the interaction of a neutral beam (with velocity \mathbf{v}_b) with the local magnetic field \mathbf{B} through $\mathbf{E} = \mathbf{v}_b \times \mathbf{B}$ [Hawk 99, Levi 90]. For different lines-of-sight, the polarization angle of the Doppler shifted Stark D_α emission lines, linearly polarized parallel (π) or perpendicular (σ) to the electric field, yields the direction of the local magnetic field [Levi 89, Levi 90].

Using a dual modulator polarimeter apparatus employing two birefringent Photo-Elastic Modulators (PEMs), running at $f_1 \approx 20$ kHz and $f_2 \approx 23$ kHz and modulating the light polarization angle about its nominal value, a linear polarizer and an Avalanche Photo-Diode (APD) sensor, the average local light polarization angle γ due to the Lorentzian electric field becomes encoded as a combination of AM components in an electric signal [Hawk 99, Stra 99], see figure 3.2. At JET, 25 channels covering

¹By *bursty* one means frequently interrupted, transient or pulsating.

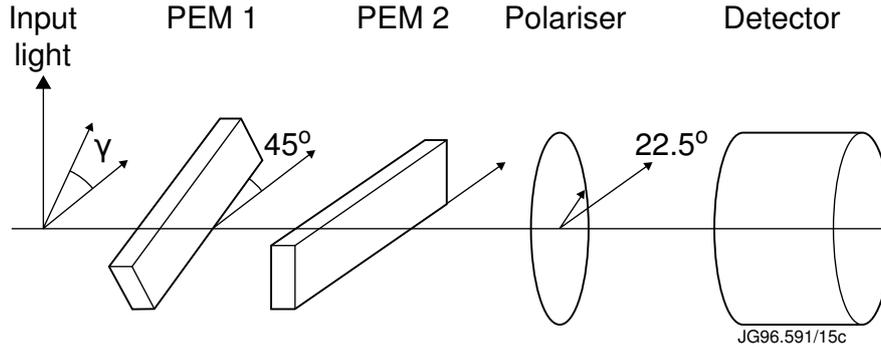


Figure 3.2: Simplified single channel light path for the MSE diagnostic at JET. *Courtesy of EFDA-JET.*

the low-field-side with ≈ 5 cm resolution employ interference filters rejecting all light but the Doppler shifted π^+ D_α line and γ is calculated through (3.12), where A_x denotes the non-stationary amplitude of APD signal component x and C_{yz} is a calibration coefficient [Hawk 06].

$$\tan [2\gamma(t)] = \frac{C_{12}A_{DC}(t) + C_{22}A_{f_2}(t) + C_{23}A_{2f_2}(t) + C_{24}A_{2f_1}(t)}{C_{11}A_{DC}(t) + C_{12}A_{f_2}(t) + C_{13}A_{2f_2}(t) + C_{14}A_{2f_1}(t)} \quad (3.12)$$

The APD signal itself is a complex spectral combinatory mix of non-stationary PEMs-related components, their higher-order harmonics and all possible inter-component beatings [Shi 06] plus a noise term $\epsilon(t)$ yielding equation (3.13).

$$S_{APD}(t) = \sum_i A_i(t) \times \cos [\omega_i(f_1, f_2)t + \phi_i] + \epsilon(t) \quad (3.13)$$

In order to extract the amplitudes of the required signal components directly related with the PEMs' excitation waveforms from (3.13) and use them in (3.12) to calculate the polarization angle, conventional lock-in amplifier schemes are typically used [Levi 89, Hawk 06, Alve 04, Alve 06, Shi 06].

Although simple in its (Fourier) essence and implementation, a SPL is naturally sensitive to offset phase mismatches between the PEMs' excitation waveforms and the relevant components in the APD signal. It is straightforward to verify that a phase mismatch of $\approx 18^\circ$ in the mixed signals will propagate into an artificial 5% decrease in the amplitude calculation and ultimately with deleterious consequences in the pitch angle calculation. Dual Phase Lock-in (DPL) implementations, although immune to this issue since in-phase and quadrature projections are simultaneously performed, are particularly vulnerable to noise induced root-mean-square (rms) errors in the amplitude calculation and an inherent difficulty in distinguishing amplitude sign changes, especially relevant whenever the pitch angle changes sign². Furthermore, being founded in Fourier analysis and therefore operating on windowed data blocks, these methods are inevitably a delicate compromise between noise attenuation and output rate in the context of a multi-rate signal processing scheme.

In addition to these difficulties, bursty events such as ELMs perturb significantly the MSE measurements [Alve 06, Jaku 07] by greatly enhancing the background radiation thus leading to implausible pitch angle measurements. Yet another inevitable source of

²Further insight on this issue will be given later on.

biasing is the interaction of the diagnostic with the power supplies of both the PEMs and the neutral beam injectors (NBIs) introducing a 50 Hz modulation in the measurements which may be non-negligible.

In an attempt to ameliorate the aforementioned issues, a KF-based technique has been developed for the real-time estimation of the amplitudes of the required components immerse in the APD signals. Provision has also been made to mitigate the biasing effect of both ELMs and the power grid 50 Hz component.

Kalman filtering model of a lock-in amplifier

Building on the basic principles of linear harmonic estimation exposed in section 3.2, one will use a multi-component lock-in amplifier implementation of a KF [Caci 96, Bitm 86, Grub 94]. The mathematical formulation of such a simultaneous multi-component harmonic estimator is given by combining the generic KF equations (2.31)-(2.35) with the harmonic estimator's state transition model given by (3.7) and (3.8) along with the predicted measurement model given by (3.9).

The lock-in analysis must be performed for each MSE channel/line-of-sight (los) for which a single measurement z_k (APD signal) is available hence the measurement noise covariance matrix R is simply a scalar quantity. The process noise covariance matrix is given by $\mathbf{Q} = q\mathbf{I}$ where q is a scalar and \mathbf{I} is the $2N_h \times 2N_h$ identity matrix thus implicitly assuming the equivalence of the uncertainty in the underlying linear process for all harmonics. Finally, using the simultaneous estimations obtained with the KFHE for the in-phase $x_{2i-1,k}$ and quadrature $x_{2i,k}$ of each signal component i , the amplitudes are obtained via (3.10).

It is clear from (3.12) what signal component amplitudes are relevant in the pitch angle calculation. However, as noted in [Levi 90] and previously mentioned, the APD signal is spectrum-wise extremely rich including not only the higher-order harmonics of the PEMs' excitation frequencies ($f_1 \approx 20$ kHz and $f_2 \approx 23$ kHz) but also all possible inter-beatings between them and their harmonics. In the context of a real-time lock-in implementation using the KFHE it is therefore crucial to estimate all relevant signal components simultaneously. Failing to do so will mean that all components with non-negligible amplitude not being estimated will (through (2.34)) contaminate the estimates of the other ones.

Another critical issue is the fact that, using (3.5), only the absolute value of the amplitude can be inferred. However, the amplitude of some harmonics can change sign. This is evident from figure 3.3(d) where the post-pulse reconstructed value of the vertical field B_z on the magnetic axis is plotted as a function of time and radial position. It is interesting to note that $B_z = 0$ along the black curve that represents the radial position of the magnetic axis. This is in fact compatible with the analysis of (3.12) for which a change in the sign of γ can only occur if one/some of the amplitudes themselves change sign. Given this situation, an approach to address this shortcoming must be devised. Considering the square reference/excitation waveform $R_f(t)$ of one of the PEMs³, with fundamental frequency f , and the single-tone KFHE defined by (3.3) and (3.4) where $\Delta\phi_k = 2\pi f/F_s$, one can estimate the in-phase $R_{f,c}(t)$ and quadrature $R_{f,s}(t)$ sinusoidal reference waveforms and compare those with the in-phase and quadrature components of the APD signal's estimated state variables $x_{f,c}(t)$ and $x_{f,s}(t)$. Using these, the absolute phase difference between the modulation in the reference and the APD measurement $\Delta\Psi$ is fully defined by (3.14).

³At JET the PEM's references are zero-mean square waveforms.

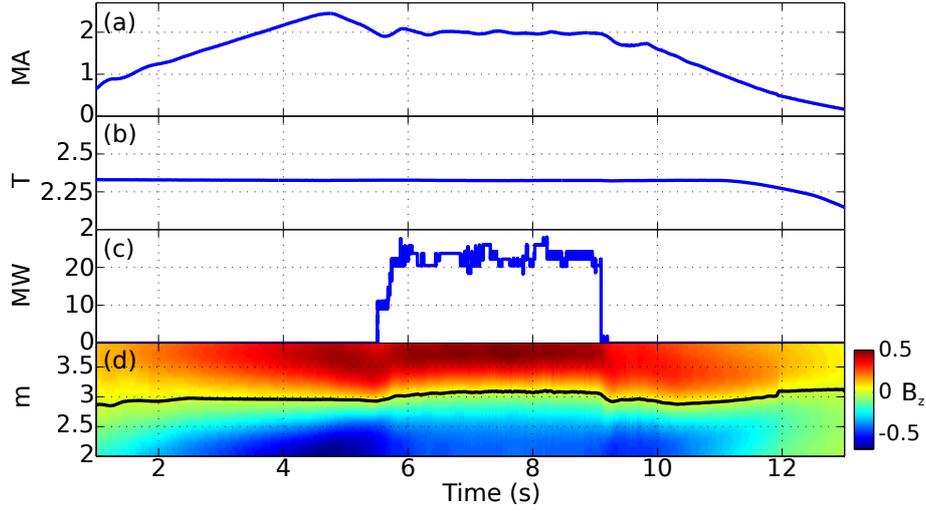


Figure 3.3: JET pulse #85394 - (a) Plasma current; (b) Toroidal field; (c) NBI power; (d) Vertical field B_z and radial position of the magnetic axis (black curve) as calculated by an off-line equilibrium reconstruction code.

$$\begin{cases} \cos(\Delta\Psi(t)) = x_{f-c}(t)R_{f-c}(t) + x_{f-s}(t)R_{f-s}(t) \\ \sin(\Delta\Psi(t)) = x_{f-s}(t)R_{f-c}(t) - x_{f-c}(t)R_{f-s}(t) \end{cases} \quad (3.14)$$

Whenever there is a sign change in the amplitude of components related to f , there will also be a sign change in, at least, one of the previous expressions.

Fourier-based lock-in methods

Standard Fourier-based lock-in amplifier schemes rely on the orthogonality of sinusoidal functions when mixing the detected signal with a reference signal. They are basically homodyne detectors in tandem with a typically narrow band Low-Pass Filter (LPF). Particularising for the case of the MSE diagnostic at JET, let one consider the unitary amplitude version of one of the PEMs' excitation waveforms $R_c(t)$ and its quadrature $R_s(t)$ given by (3.15), where ϕ is a tunable phase offset.

$$R_c(t) = \cos(\omega_{PEM}t + \phi) \quad (3.15)$$

$$R_s(t) = \sin(\omega_{PEM}t + \phi)$$

For simplicity and without loss of generality, let one assume an APD signal of the form $S_{APD}(t) = a(\omega_{PEM}, t) \cos(\omega_{PEM}t) + \epsilon(t)$ where $\epsilon(t)$ is a white noise term. The DC component of the mix of these references with the APD signal is directly proportional to the amplitude of the component of interest as is clear from (3.16). These are simply the cosine and sine transforms inherent to the discrete Fourier transform algorithm.

$$\langle S_{APD}(t) \times R_c(t) \rangle_{LPF} = \frac{1}{2}a(\omega_{PEM}, t) \cos(\phi) \quad (3.16)$$

$$\langle S_{APD}(t) \times R_s(t) \rangle_{LPF} = \frac{1}{2}a(\omega_{PEM}, t) \sin(\phi)$$

Compensating for the delay between the reference and the APD signal at frequency ω_{PEM}

implies that $\phi = 0$ and, consequently (3.17). This is the SPL.

$$a(\omega_{PEM}, t) = 2 \times \langle S_{APD}(t) \times R_c(t) \rangle_{LPF} \quad (3.17)$$

In case delay compensation is not straightforward, $|a(\omega_{PEM}, t)|$ can still be recovered from (3.18) through a DPL.

$$|a(\omega_{PEM}, t)| = \sqrt{[2 \times \langle S_{APD}(t) \times R_c(t) \rangle_{LPF}]^2 + [2 \times \langle S_{APD}(t) \times R_s(t) \rangle_{LPF}]^2} \quad (3.18)$$

Although the DPL does not require phase compensation like the SPL does, it is incapable of detecting sign changes in the amplitude (unavoidable when considering the polarization angle measured by two channels/lines-of-sight on opposite sides of the plasma's magnetic axis). Furthermore the DPL, unlike the SPL, is vulnerable to root-mean-square (rms) errors simply because $|a(\omega_{PEM}, t)| > 0$, even if the true amplitude is zero, just because there is a finite noise variance ($\sigma^2[\epsilon(t)] \neq 0$) in the APD signal.

Digital implementations of Fourier-based lock-in schemes require the windowed calculation of one or both expressions in (3.16) depending on the chosen lock-in implementation. The narrower the LPF bandwidth, the better the Signal to Noise Ratio (SNR) in the amplitude estimation at the expense of a higher latency.

Amplitude frequency response

The performance of the KFHE, as with all classical filters, is basically a balance between SNR and delay/reaction to transients. Typically the better the SNR the greater the latency in the estimates and vice-versa [Caci 96, Bitm 86, Grub 94].

In the context of a real-time signal processing scheme, the Fourier-based lock-in amplifier methods take a windowed (sample block) fragment of the measured signal and subsequently mix it with the sinusoidal reference(s). On the contrary, KF-based methods are sample-by-sample oriented rather than sample-block oriented. This means that, if one wishes to compare the spectral characteristics of both methods one needs to, somehow, find a common ground. The idea is therefore to investigate the spectral content of a pure DC signal after the windowing takes place (part of the Fourier-based lock-in algorithm), and compare it to the spectral content, again after the same windowing procedure takes place, of the KFHE's in-phase estimations of a range of frequencies. It is well known, from basic signal processing theory, that windowing methods, despite their benefits e.g. in terms of side-lobe attenuation, affect the true spectrum of a signal. In a sample-block oriented method this process is therefore inevitable so, for this study one has used a window of 15 samples. Figure 3.4 illustrates the different amplitude responses obtained with and without applying the KFHE prior to windowing. The solid black line (without symbol) in figure 3.4 shows the amplitude spectral content of a 15 sample unitary signal. It is basically the response of a 15 sample square window. Zero-padding was used in order to interpolate the results and increase the apparent frequency resolution. As mentioned before, this sort of windowed/sample-block processing is particularly relevant in the context of real-time digital implementations of Fourier-based lock-in amplifiers and, for this case in particular, it would mean that an amplitude *estimation* would be provided every 15 samples.

Regarding the KF-based method, several single-tone KFHEs are executed in parallel and each one of those is instructed to estimate a frequency which is part of the grid obtained in the previous method using the same unitary signal as a measurement. The

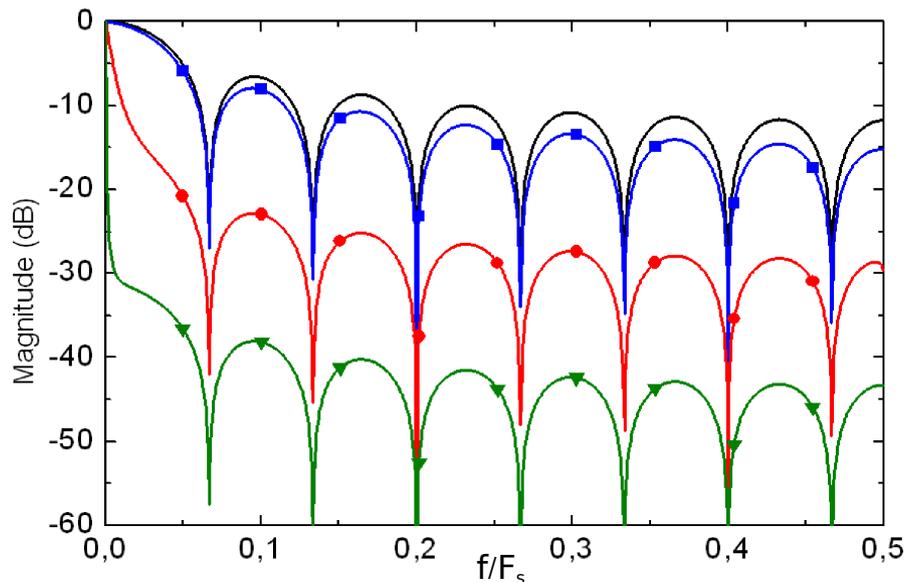


Figure 3.4: Frequency response of a 15 sample square window using the DFT alone (black line) and jointly with the KF, where the r/q ratios are equal to 10^{-4} (blue), 10^1 (red), and 10^2 (green).

various KFHEs are basically band-pass filters applied to a unitary signal and this process is therefore performed on a sample-by-sample basis. The in-phase components of each of these different frequency estimations is then subject to the exact same windowing procedure so that an equal ground comparison can take place. Finally, the spectral content of each of the in-phase estimates (at each particular frequency) is obtained using the Discrete Fourier Transform (DFT). For different values of $\lambda = r/q$, the KFHE results are shown by the coloured lines with symbols in figure 3.4. It is clear that the KFHE can provide a SNR improvement because, unlike in the previous method, before the window is applied there is actually some filtering taking place. This effect is ever more evident as λ increases making the KF privilege the model rather than the measurements. On the other hand, the lower the λ , the closer the KFHE gets to the Finite Impulse Response (FIR) filter.

It is worth noting that the performance differences observed between the Fourier-based and KF-based implementations are particularly relevant in the cases where only a limited amount of samples are available. Obviously, the phase response will complement this study as it provides information on estimation latencies. In the case of the simple average LPF used in the frequency response assessment, being a symmetric FIR, it exhibits a linear phase response, i.e. frequency independent constant delay equal to half of the number of coefficients. The KF lock-in amplifier's phase will be addressed next.

Phase response of the amplitude estimation

The comparative study on the frequency response of the amplitude estimations performed by the Fourier-based and KF-based methods provides a good indication of the potential superiority of the KF in this respect. This information obviously needs to be complemented by a study of the phase response which, in broad terms, determines the delay or the reaction time of the method to transient events. In typical large tokamaks, such as JET, the time scales at which the magnetic field can change may be as fast as < 1 ms,

in the case of MHD instabilities, and as slow as > 100 ms in the case of the global plasma equilibrium. Clearly, in order to obtain the required resolution in the pitch angle calculation, the amplitude estimations themselves are required not only to be accurate and to maximize the SNR, but also to be performed with the smallest possible latency. This would be even more crucial when devising feedback control schemes addressing, for instance, MHD control.

Let one consider a simple, single carrier, AM signal immerse in additive white noise and given by (3.19), where ω_c is the known carrier frequency, ω_M is the unknown modulation frequency, $v(t)$ is the noise term and A and A_M are, respectively, the carrier and modulation amplitudes.

$$S(t) = (A + A_M \cos(\omega_M t)) \cos(\omega_c t) + v(t) \quad (3.19)$$

By rewriting the previous as (3.20), as expected, it is clear that any AM signal can also be interpreted as the sum of three components typically know as the carrier ω_c and the sidebands $\omega_c \pm \omega_M$.

$$S(t) = A \cos(\omega_c t) + \frac{A_M}{2} \cos((\omega_c - \omega_M)t) + \frac{A_M}{2} \cos((\omega_c + \omega_M)t) + v(t) \quad (3.20)$$

In a single-tone implementation of a KF-based lock-in amplifier, (3.5) is therefore expected to provide the total carrier amplitude $A_c = A + A_M \cos(\omega_M t)$. However, as discussed previously, in a single-tone implementation of the KFHE where ω_c is being estimated, the sidebands themselves $\omega_c \pm \omega_M$ will contaminate the estimates as they are not accounted for in the KF's state transition model. So, although the in-phase estimation of the ω_c component should indeed be in-phase with the true component, the same may not hold for the sidebands. Furthermore, because the KFHE will attenuate all components not accounted for in the state transition model, the amplitude of the sidebands may be, depending on λ , smaller than $A_M/2$.

With this in mind, two studies were performed on the frequency response of the KFHE in the case where not only the carrier ω_c is present, but also an intruder component mimicking the effect of one of the AM sidebands is present. On the first study, the KFHE is instructed to estimate the DC component of a signal which contains not only DC, but also an *intruder* component. The sampling frequency F_s is 250 kHz. A set of *intruder* components is defined in the range $0 - F_N$ where $F_N = F_s/2$ is the Nyquist frequency. For each of these *intruders*, a measurement signal is simulated and the KFHE is subsequently instructed to estimate its DC component. Finally the DFT is performed on the KFHE's in-phase estimate and its complex value, taken at the value of the *intruder* frequency component, is stored. The result for the amplitude and phase's frequency response is shown, respectively, in figure 3.5(a) and (b). The various coloured curves correspond to distinct values of r having $q = 10^{-4}$. As expected, the larger the λ i.e. the more the KF believes in the state transition model, the greater the attenuation of the *intruder* frequency component and the larger is its delay when contaminating the in-phase estimation of the DC component. Similarly, the lower the λ , the lower the attenuation of the *intruder* frequency component and the lower is its delay when contaminating the in-phase estimation of the DC component. It is interesting to observe that the delay never exceeds a quarter of each component's period.

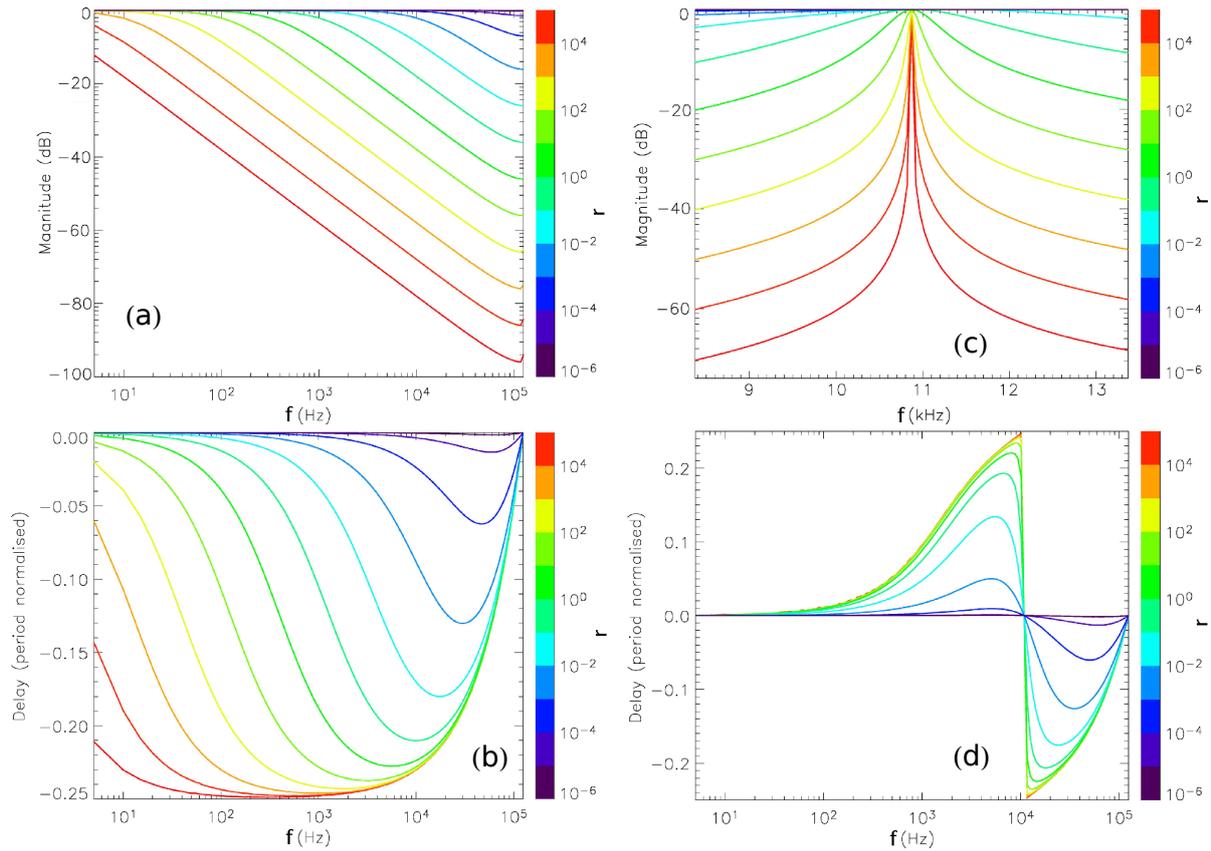


Figure 3.5: (a) Amplitude and (b) phase response of the low-pass estimator; (c) Amplitude and (d) phase response of the band-pass estimator.

On the second study, the KFHE is instructed to estimate the 10.87 kHz⁴ component of a signal which contains not only that same frequency, but also an *intruder* component. Again, a set of *intruder* components is defined in the range $0 - F_N$ where $F_N = F_s/2$ is the Nyquist frequency and, for each of these scenarios, a KFHE is performed and the contamination of that spurious component in the in-phase estimate is characterised. The results are shown, in figures 3.5(c) and (d). Again, it can be seen that the phase mismatch never exceeds a quarter of the *intruder's* period and that, the larger the attenuation, the larger the delay.

In general, and unlike a symmetric FIR filter, the KFHE's phase response is not linear. Typically, a reasonable balance between SNR and latency is achievable for values of r in the range $10^0 - 10^1$. Regarding the band-pass frequency estimation, figure 3.6 suggests that the normalized delay is approximately symmetric about the estimated component, assuming that the modulation frequency is much smaller than the carrier frequency. Using this observation one can write (3.21) for the in-phase estimation $S_{est,c}$ of the frequency component ω_c , where A_{MRl} and A_{MRr} are, respectively, the left and right sidebands' residual amplitudes ($A_{MRl} \approx A_{MRr} \approx A_{MR}$), ϕ is the phase lag of the estimated sideband components and $v_R(t)$ is the residual noise. Basically the subscript R denotes the residual of a quantity that persists after the KFHE. It is worth noting that A_{MR} may be lower than A_M because only ω_c is part of the KFHE's state transition model.

⁴Chosen so that it wouldn't be a sub-multiple of the sampling frequency hence unprivileged with respect to any other frequency (lower than Nyquist's) in the real axis.

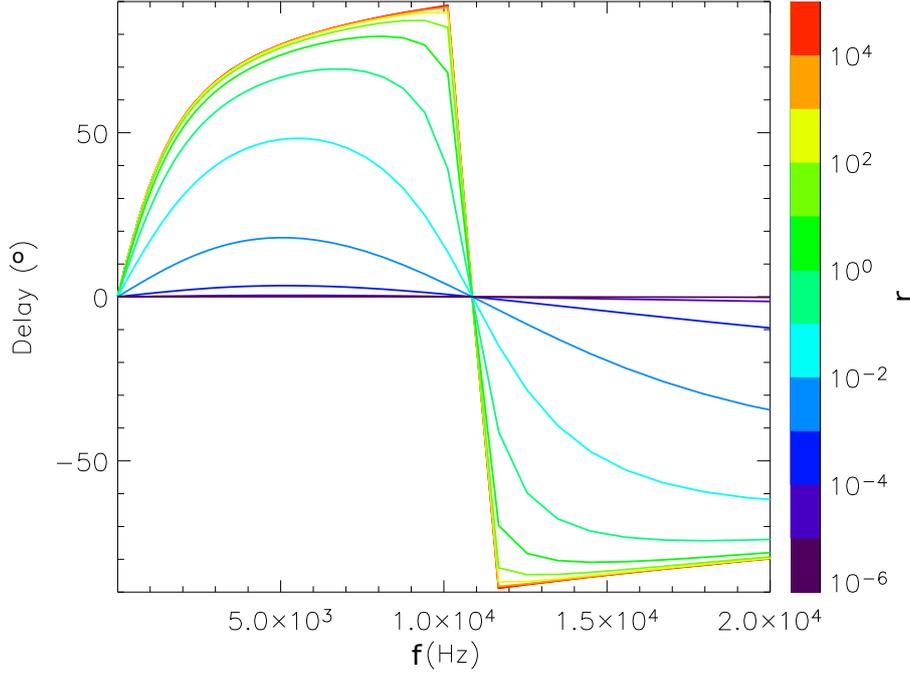


Figure 3.6: Phase response zoom of the band-pass estimator.

$$S_{est.c}(t) = A \cos(\omega_c t) + \frac{A_{MRl}}{2} \cos((\omega_c - \omega_M)t + \phi) + \frac{A_{MRr}}{2} \cos((\omega_c + \omega_M)t - \phi) + v_R(t) \quad (3.21)$$

After basic manipulation, the previous can also be written as (3.22) which, as expected, accounts for a delay in the estimation of the modulation amplitude. This delay happens to be the same as the delay in the sideband estimation which, obviously depending on λ , is bounded by a quarter of a period.

$$S_{est.c}(t) = (A + A_{MR} \cos(\omega_M t - \phi)) \cos(\omega_c t) + v_R(t) \quad (3.22)$$

Frequency mismatches

Regardless of whether a Fourier-based or KF-based lock-in technique is used, ω_c is assumed to be known. However, taking the case of the MSE diagnostic as an example, frequency jitter (e.g. due to thermal effects) in one of the PEMs' reference waveforms will inevitably impoverish the quality of the amplitude estimations. For example in the case of the SPL, a phase mismatch of $\sim 25^\circ$ accounts for a reduction of $\sim 10\%$ in the estimated amplitude.

In order to test the KF's method sensitivity to mismatches between the carrier frequency used in the state transition model and the true PEMs reference frequency, a simulated signal $S(t) = A(t) \cos(\omega_c t + \phi_j(t)) + v(t)$ where $\omega_c/2\pi = 20$ kHz, $A(t) = 10 + 3 \cos(\omega_M t)$, $\omega_M/2\pi = 20$ Hz, the unknown phase $\phi_j(t)$ is given by (3.23) and $v(t)$ is a white noise term with $\sigma[v(t)] = 4$ V, was built.

$$\begin{cases} \phi_j(t) = 0 & , 0 \leq t < 0.2 \text{ s} \\ \phi_j(t) = \Delta\omega_j t & , 0.2 \leq t \leq 0.4 \text{ s} \end{cases} \quad (3.23)$$

Three cases are tested: $\Delta\omega_j/2\pi = 10^1, 10^2$ and 10^3 Hz. Figure 3.7(b) shows an example of the amplitude estimated with the KF-based lock-in implementation (orange curve) and

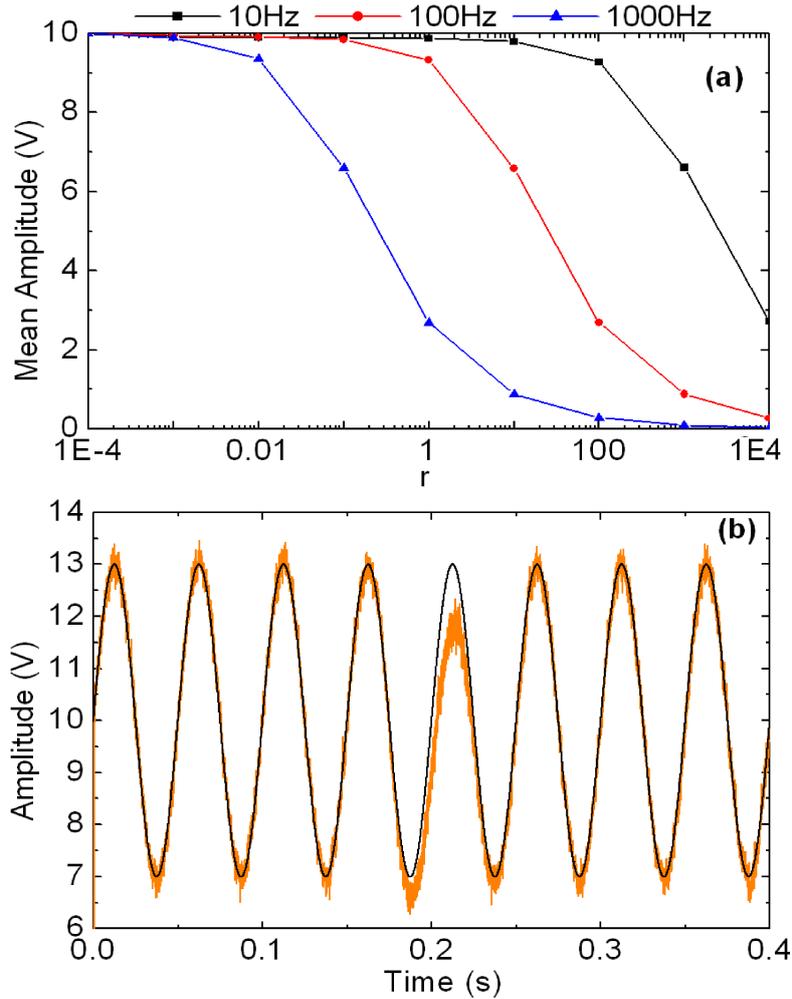


Figure 3.7: Effect of a frequency mismatch on the estimated amplitude of an AM signal ($q = 10^{-4}$). (a) Offset in the mean amplitude estimation increases with covariance r and mismatch frequency $\Delta\omega/2\pi$. (b) (Orange) Time evolution of the estimated amplitude for the $r = 1$ and 100 Hz mismatch case is shown.

the true modulation (black curve) for $r = 1$ and $\omega_j/2\pi = 10^2$ Hz. In fact, throughout these tests a $q = 10^{-4}$ was used. It is evident that the frequency mismatch produces a transient loss in the amplitude tracking which is subsequently recovered. Figure 3.7(a) summarizes the results of the test where, for $r = 10^{-4} - 10^4$, the mean value of the amplitude estimation is shown. While a frequency mismatch of 1 kHz is tolerable down to $r = 10^{-3}$, a mismatch of 10 Hz is still tolerable for $r = 10^1$. Basically, the larger the frequency mismatch, the more the state transition model diverges from the underlying process. Still, the lower the r , the more the KF trusts the measurements and less the model hence it is natural that, for sufficiently low values of r , even large frequency mismatches are tolerable.

MSE signal processing at JET

The KF-based lock-in amplifier implementation has been shown to be a suitable candidate for the real-time estimation of the required frequency component amplitudes that, through (3.12), enable the pitch angle calculation. The SPL method is adequate for post-

pulse analysis, where narrow-band low-pass filters (eventually without delay since both the past and the future of the measurements are known offline) can be used and the phase difference between the reference sinusoidal waveforms and their counterparts in the APD signal is known with sufficiently good accuracy. However, on a real-time implementation, the Fourier nature of the method restricts its performance. Unlike the KF-based method, the SPL has no filtering capability prior to the windowing process hence, not only it strives to remove noise but it also exhibits a poor resilience to the contamination of neighbouring components. This effect becomes more pronounced as the number of samples in the windowing process becomes small. Indeed the SPL implements a LPF in its design, but it is only applied after the windowing takes place thus after the spectral contamination. In contrast with the SPL method, which gives an average value (over the last few samples determined by the box window) of the amplitude estimates, the KF-based method works on a sample-by-sample basis. As mentioned previously, phase mismatches between the reference and the actual frequency component in the APD signal will inevitably result in an underestimation of the amplitudes. The DPL and the KF-based methods are immune to this problem as they perform the signal projection on the complete basis (*cosine* and *sine*). However, unlike the KF-based method and the DPL, the SPL does not provide biased amplitude estimates and is inherently capable of taking into account sign changes in the amplitude estimations at the expense of knowing in advance the phase difference between the reference waveform and its corresponding frequency in the APD signal. Still, as addressed later on, it is possible to circumvent this apparent limitation of the KF-based method at the expense of an extra processing demand.

In parallel with these signal processing issues, at JET further difficulties arise. Plasma perturbations, in the form of type-I ELMs, significantly affect the pitch angle estimation. The physical mechanism for this occurrence is not yet completely understood but it may be related to the substantial increase in the background radiation so, it is clear that these abrupt changes in the pitch angle (hence also in the q-profile), measurable even by MSE channels probing the plasma core, are implausible.

Building on the previous assessments, a KF-based lock-in implementation has been developed, as a proof of principle, to target the real-time signal processing of the MSE data at JET. The state transition model, measurement and amplitude estimation expressions are based on (3.7)-(3.10). In fact, at JET, a multi-tone KFHE is used at the core of the KF-based lock-in amplifier scheme. In this scheme, ~ 15 components (including the ones used in (3.12) and other dominant beating combinations) have been identified as relevant in the APD signals and are therefore simultaneously estimated. The total spectrum of the APD signal for MSE channel 15 highlighting the components being simultaneously estimated is depicted in figure 3.8. Although this scheme does provide real-time estimates for the amplitudes of each harmonic with good SNR, it has some limitations. The first is related to the sign of the amplitude estimation given by (3.10). As mentioned previously it might occur that, for the particular diagnostic's viewing geometry and in particular for channels collecting light from opposite sides of the magnetic axis, some component amplitudes have opposite sign. In fact, this is the typical behaviour of the 23 and 40 kHz harmonics at JET. As explained in section 3.3.1, having a KFHE tracking the PEMS' reference signal in parallel with the KF-based lock-in implementation working on the APD signal can, through (3.14), provide the means to determine the correct sign and an eventual flipping due to modifications in the magnetic equilibria.

Given that any relevant component not included in the state transition model will inevitably propagate into the other components' estimations, it may be convenient to

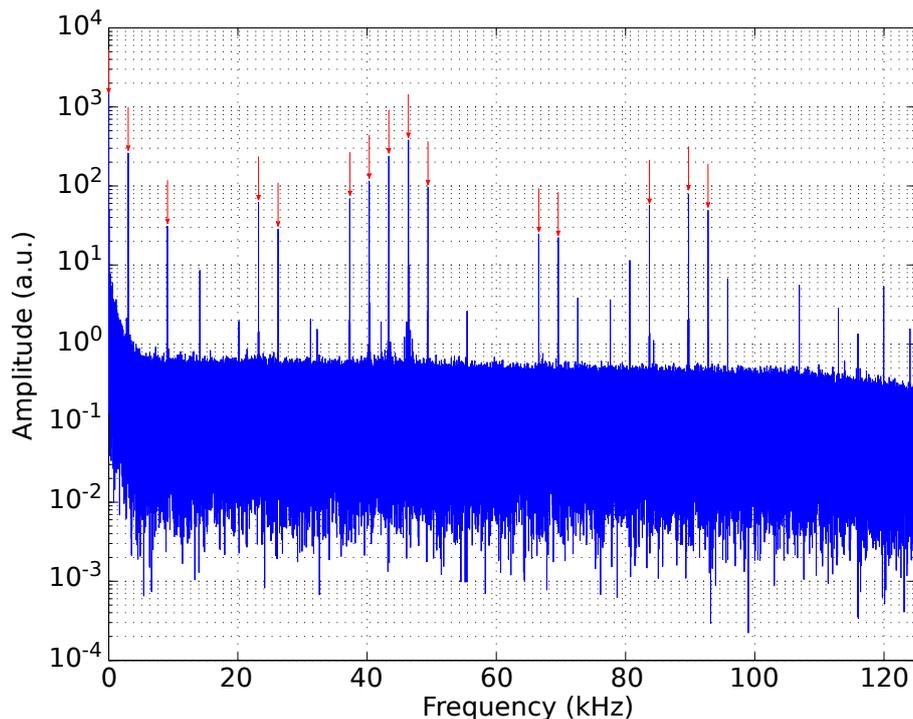


Figure 3.8: Spectrum of channel 15's APD signal for JET pulse #70129 indicating which of the frequency tones are used in KFHE's state transition model.

take into account the grid induced 50 Hz so as to minimize the contamination in the pitch angle calculation. In this case, two approaches emerge. Either the 50 Hz estimation is an integral part of the simultaneous amplitude estimation of the various components, or it is performed after the previous but only on the components used in (3.12). The latter is clearly the most rigorous in mathematical/scientific terms however, it does approximately double the processing burden if one recognizes that, for each relevant harmonic, two 50 Hz sidebands exist corresponding to the amplitude modulation. Not only the extra processing burden is intolerable, the benefits in terms of results do not justify it. So, the approach taken is to apply a second filter bank with a smaller $\lambda = r/q$ (for a fast time response) to the streaming amplitude estimates, tracking only DC and 50 Hz. The DC component gives the desired result.

To what concerns the implausible pitch angle calculations at the time of bursty ELM activity, it is clear that a strategy for mitigating this effect is certainly desirable. It is also clear from figure 3.9 the notorious correlation between the peaks in the D_α for an outer divertor line-of-sight and the APD signal from channel 15 (major radius $R \sim 3.225$ m) at the time of an ELM. The gray curve in the same figure represents the estimated DC component. As concluded in section 3.3.1, a good compromise between SNR and estimation latency is obtained with $\lambda = 10^4 - 10^5$ in normal operation conditions, i.e. *outside* ELMs. The idea is to change the gain of the multi-tone KFHE according to (3.24) i.e. whenever the APD signal z_k is above a percentage δ of the KFHE's DC in-phase estimation $x_{DC-,k}$, the KF's gain K_k is reduced by three orders of magnitude with respect to its original value K_o immediately before the signal rise due to the ELM. Similarly, once the APD signal decays to a value which is back to a closer vicinity of $x_{DC-,k}$, the original

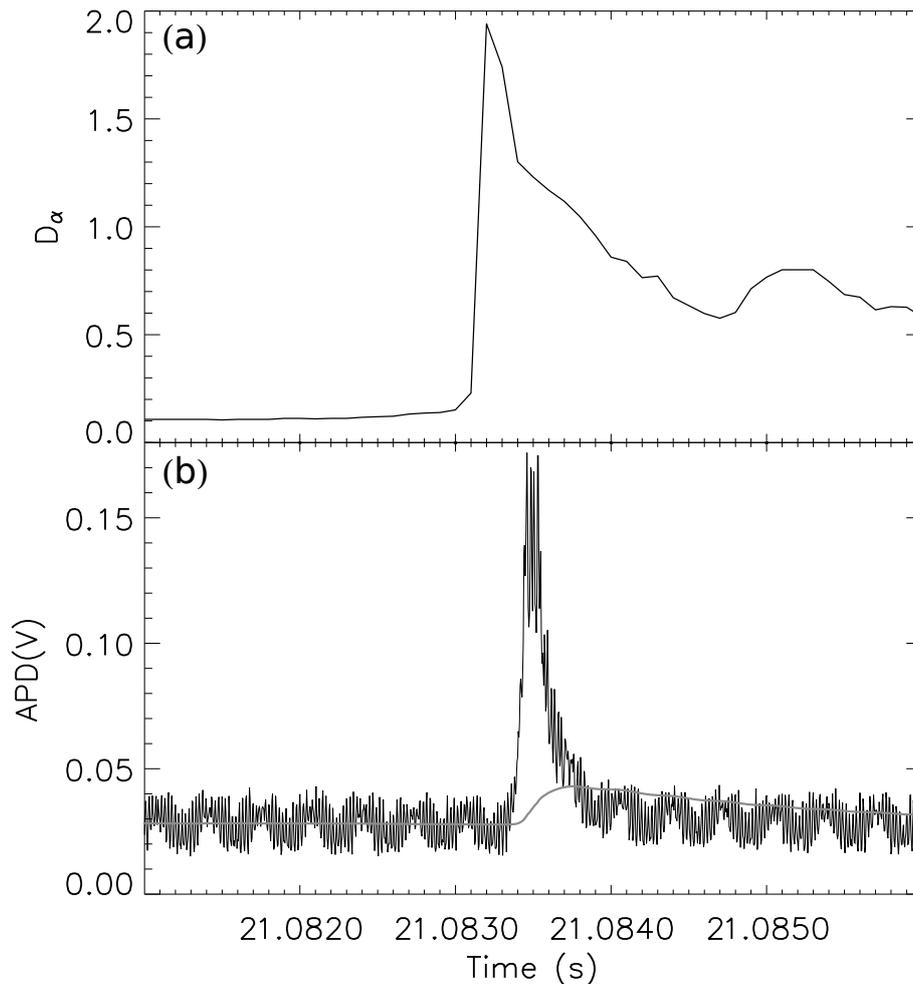


Figure 3.9: (a) D_α radiated power at the outer divertor leg; (b) APD signal of channel 15 (black curve) and DC component estimation of the multi-tone KFHE (gray curve).

gain before the ELM is restored and normal operation is resumed.

$$\begin{cases} K_k = K_o \times 10^{-3} & , (z_k - x_{DC,c,k})/x_{DC,c,k} > \delta \\ K_k = K_o & , (z_k - x_{DC,c,k})/x_{DC,c,k} \leq \delta \end{cases} \quad (3.24)$$

Using this scheme and whenever an ELM is detected the KF gain is substantially reduced so that the APD measurements are essentially ignored causing the KFHE to trust, almost exclusively, the state transition model. Likewise, whenever the ELM's effect on the APD wears off, the filter resumes its normal operation and the original gain is restored. In physical terms, the approach is to consider that during an ELM there should be little or no variation of the plasma's safety factor profile.

Experimental results

The KF-based lock-in amplifier scheme featuring 50 Hz and ELM suppression has been used to provide pitch angle measurements in JET pulse number 70129. The analysis was performed offline but the implementation itself was performed as if it was processing data in real-time. The pulse itself is a 2.3 MA/1.76 T with 10 MW of NBI power and a (type-I) ELM frequency of the order of 15 Hz.

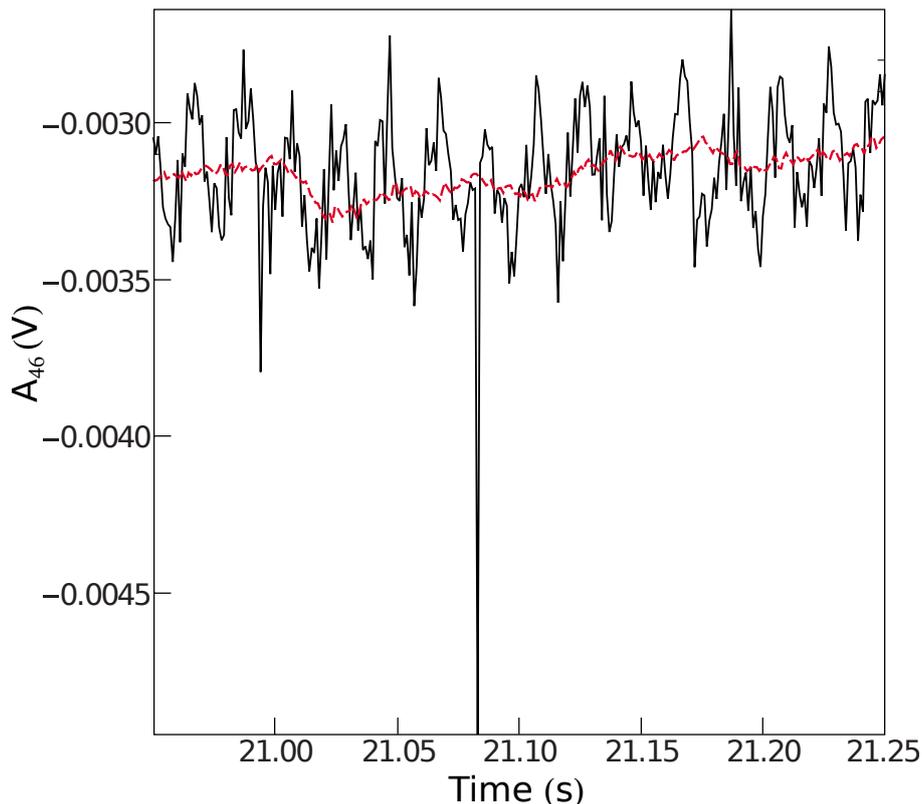


Figure 3.10: Amplitude of the 46 kHz component using a lock-in amplifier (black) and the Kalman filter algorithm (red) with ELM+50 Hz removal. The ratio r/q , set to 10^5 in between ELMs, is reduced by three orders of magnitude during the ELMs

At JET the raw APD data for each of the 25 channels plus the PEMs' excitation waveforms is sampled at a rate of 250 kSample/s. Results obtained with the present SPL-based real-time system are also used for comparison purposes. The latter runs on a 1 ms cycle thus processing blocks of 250 samples on each cycle. A Hanning window is used and subsequently a box-car average acts as the SPL's LPF. This mechanism provides a value for each of the required amplitudes and pitch angle γ , for each line-of-sight, every millisecond. Finally the pitch angle undergoes an additional filtering stage (third order, elliptical, Infinite Impulse Response (IIR) LPF with a cutoff at 35 Hz) to compensate for the poor SNR provided by the SPL [Alve 04]. In order to perform a common ground comparison between both methods, the same 250 point box-car averaging is also used in the KF algorithm (although no Hanning window is used), providing results at the same rate as the SPL algorithm.

The superior performance of the KF-based lock-in is patent in figure 3.10 where the time trace of the amplitude of the 46 kHz component in the presence of an ELM event is shown. The black curve and the red curve are, respectively, the results obtained with the SPL and KF-based method. It can be seen that the SPL result is contaminated with the 50 Hz component and exhibits a strong spike when the ELM occurs. Furthermore, as is also evident from figure 3.11, the KF-based method provides amplitude estimations with a larger SNR than the SPL, even without removing the 50 Hz and mitigating the effect of the ELMs. Eventually, for a larger box-car averaging (narrower low-pass filtering), the results from both methods would become similar.

In an effort to compare the impact of the amplitude estimations, provided by the

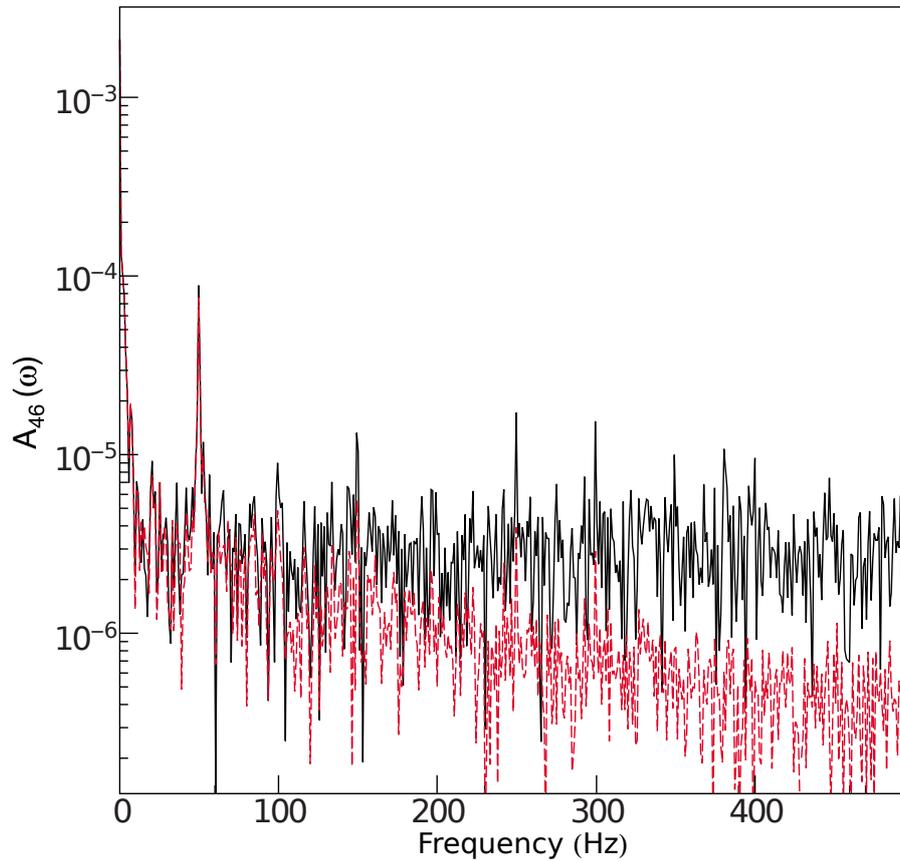


Figure 3.11: Spectrum of the 46 kHz component using a the SPL (black) and the KF-based method (red) with the ELM+50 Hz removal controller disabled

SPL and KF-based methods, in the actual pitch angle calculation, a study was performed and its results are resumed in figure 3.12. Before the ELM (figure 3.12(a)), both methods present comparable results although the radial profile of the pitch angle calculated using the SPL denotes a higher variance most likely due to the lower SNR of its amplitude estimations. Subsequently during the ELM (figure 3.12(b)), whereas the results provided by the SPL are rather implausible, the radial profile provided by the KF-based method remains undisturbed. Finally after the ELM (figure 3.12(c)), results from both methods become again comparable.

3.3.2 Coherence spectrum

In this section, a new method for the real-time mode number estimation of MHD perturbations in tokamak plasmas is described. At its core is the multi-tone KFHE already described and the idea is to use its spectral estimation capabilities to perform real-time cross spectral analysis of multi-sensor data. As an example, the method is used for the analysis of a typical ISTTOK pulse with bursty-like MHD activity.

Introduction

The successful real-time control of tokamak plasmas depends heavily on the quality of the signal processing methods employed in diagnostics and real-time processing nodes. For instance, low noise and low latency estimates of the frequency, amplitude, phase and mode

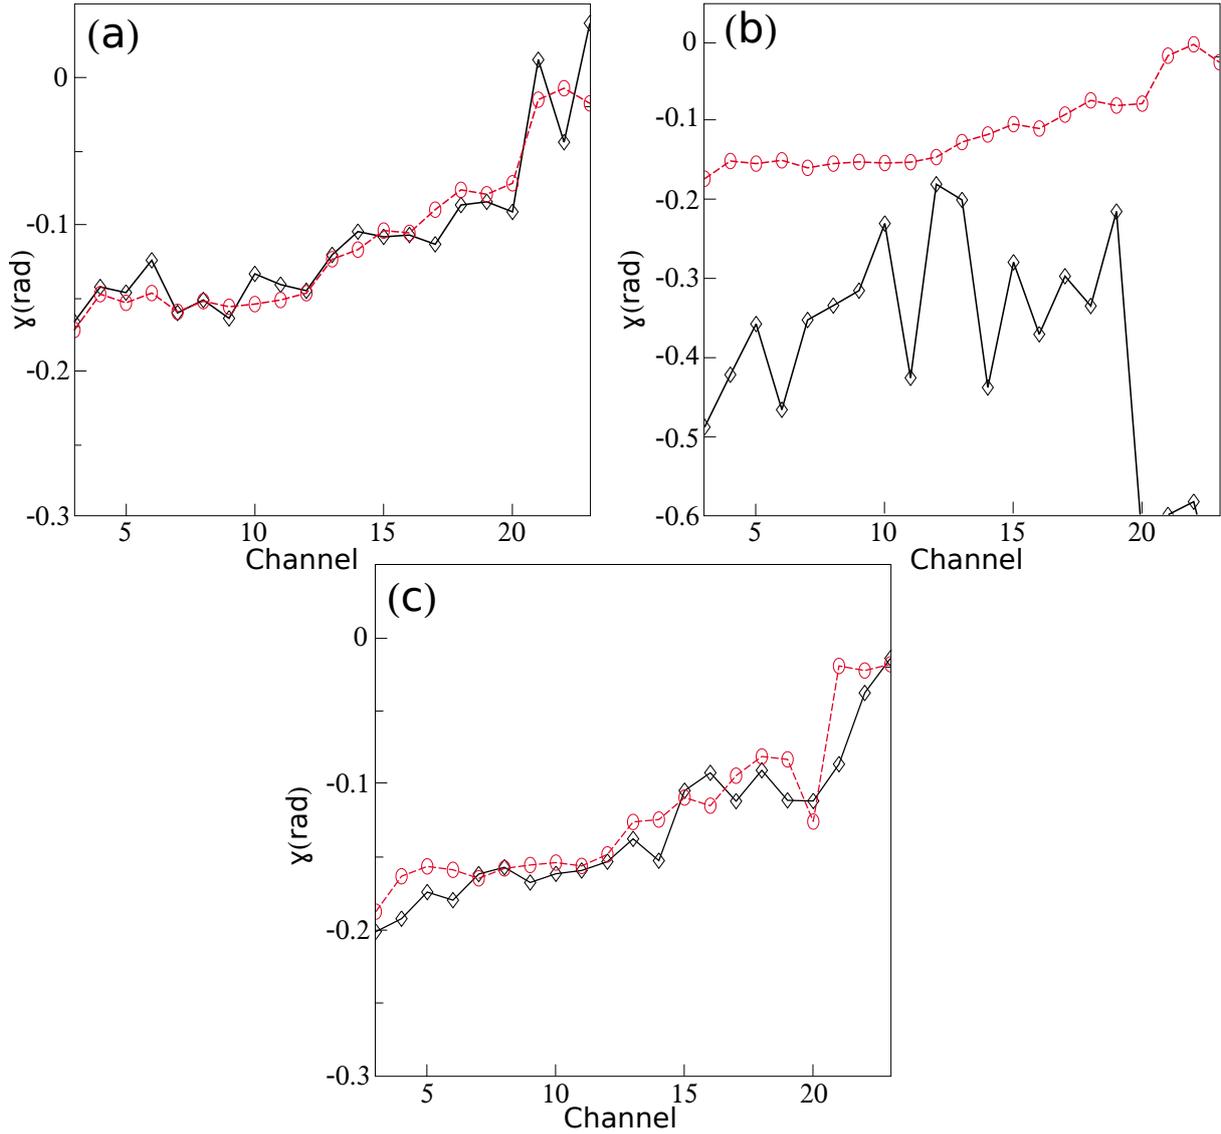


Figure 3.12: Radial profile of the pitch angle calculation using the SPL (black) and the KF-based method (gray) at: (a) $t \sim 61.0825$ s; (b) $t \sim 61.0845$ s and (c) $t \sim 61.0925$ s

number of MHD activity are invaluable for the real-time control of MHD activity such as neoclassical tearing modes (NTMs) [Berr 06] or resistive wall modes (RWMs) [Orto 06]. In the following sections, one addresses the characterization of the wavenumber spectra of unstable modes observed in ISTTOK plasmas using a set of two poloidal arrays of magnetic sensors (Mirnov coils) measuring dB_θ/dt where B_θ is the poloidal magnetic field. In particular, the prospect for the estimation of the poloidal wavenumber in real time using Kf-based techniques is investigated. It is shown that, in contrast with typical cross power spectrum analysis, where relatively large datasets are required for mode number estimation, with this innovative technique, the mode number spectra may be obtained not only in faster time scales but also with fewer samples.

Let one assume magnetic perturbations of the form $\delta B \propto \delta B(r) \cos(m\theta - 2\pi\omega t + \psi_0)$, where r is the plasma minor radius, $m(\omega)$ is the poloidal mode number, θ is the poloidal angle of a magnetic sensor, ω is the frequency at which the rotating mode is observed and ψ_0 is the initial phase of the perturbation at $t = 0$ and poloidal angle $\theta = 0$. Then, the phase difference between the signals measured by two magnetic coils (1 and 2) located at

the same toroidal position but different poloidal position is given by (3.25) where $\Delta\theta_{12} = \theta_2 - \theta_1$. This is the fundamental relation for, in this case poloidal, mode number analysis.

$$\Delta\psi_{12}(\omega) = m(\omega)\Delta\theta_{12} \quad (3.25)$$

Furthermore, let one also define the cross-power spectral density $C_{12}(\omega)$ between measurements from coils 1 and 2 as (3.26), where $S_1(\omega)$, $S_2(\omega)$, and $S_{12}(\omega)$ are, respectively, the statistically ensemble averaged auto-spectrum and the cross-spectrum densities of the two signals.

$$C_{12}(\omega) = \frac{S_{12}(\omega)}{[S_1(\omega)S_2(\omega)]^{\frac{1}{2}}} \quad (3.26)$$

From the cross-power spectral density one can calculate the *coherency* $\gamma_{12} = |C_{12}(\omega)|^2$ and phase difference $\Delta\psi_{12} = \tan^{-1} \{ \text{Im}[C_{12}(\omega)] / \text{Re}[C_{12}(\omega)] \}$. It is interesting to note that combining (3.25) and the previous yields (3.27) which provides the means to calculate the mode number m .

$$\Delta\psi_{12}(\omega) = m(\omega)\Delta\theta_{12} = \tan^{-1} \{ \text{Im}[C_{12}(\omega)] / \text{Re}[C_{12}(\omega)] \} \quad (3.27)$$

Coherent modes⁵ propagating with a given frequency ω are characterized by high coherency level ($\gamma \sim 1$ with $\gamma \in [0, 1]$), in contrast with incoherent fluctuations for which $\gamma \sim 0$. Therefore, given a reference coil and using the whole set of all other coils, it is possible to determine the coherent mode number $m(\omega)$ from the linear fit of the cross-spectrum phase $\psi_{12}(\omega)$, given by the right hand side of (3.27), as a function of the difference in poloidal angle location for each pair of coils.

The ISTTOK's turbulent nature and characteristic time scales (e.g. ~ 30 ms pulse length and $\sim 300 \mu\text{s}$ particle confinement time) imposes some constraints on the usefulness of any analysis of experimental data, namely in what concerns latency and output rate. While the Fourier-based coherence analysis is of widespread use for offline analysis of sufficiently long signals, it struggles to provide useful information whenever only a very limited amount of samples is available for analysis.

A KFHE-based coherence spectrum method has been developed with real-time capabilities in mind, but also suitable for offline analysis. The proposed method is shown to be particularly useful under the constraints of high output rates whenever only a limited amount of samples is available.

Kalman filter based coherence analysis

The same multi-component KFHE used in section 3.3.1 for the real-time signal processing of MSE data is reused here as the basis for the coherence analysis of magnetic sensor data. The idea is to use it more along the lines of [Bitm 86] by defining a linear and equally spaced frequency grid performing the simultaneous in-phase and quadrature estimation of multiple components. In order to be able to estimate the poloidal mode number of magnetic fluctuations, one has considered a hybrid scheme based on (3.26) where the Short-Time Discrete Fourier Transform (STDFT) is replaced by the multi-tone KFHE. The cross-power spectral density between the signals of two coils, labelled (1) and (2),

⁵By *coherent* one means two sensor signals whose phase difference, at the frequency at which the mode is observed, does not change over time.

is constructed from their in-phase $x_{est,c}$ and quadrature $x_{est,s}$ component estimates, for each angular frequency ω , yielding (3.28) where $\langle \cdot \rangle$ denotes temporal averaging⁶.

$$S_{12}(\omega) = \left\langle \left[x_{est,c}^{(2)}(\omega)x_{est,c}^{(1)}(\omega) + x_{est,s}^{(2)}(\omega)x_{est,s}^{(1)}(\omega) \right] + i \left[x_{est,s}^{(2)}(\omega)x_{est,c}^{(1)}(\omega) - x_{est,c}^{(2)}(\omega)x_{est,s}^{(1)}(\omega) \right] \right\rangle \quad (3.28)$$

Similar expressions for the auto-spectrum (S_1 and S_2) of each coils' signal are equivalently obtained. Having the cross and auto-spectrums, (3.26) can immediately be used to provide $\gamma_{12}(\omega)$, $\Delta\psi_{12}(\omega)$ and, finally, $m(\omega)$. Unlike the Fourier-based coherence analysis, where the total number of samples is the number of points used in each DFT times the number of time windows used in the averaging, in the case of the KFHE-based approach the minimum amount of samples is just the one required for the time average.

Because, as previously mentioned, the KFHE does provide some noise mitigation (depending on λ), the in-phase and quadrature estimates for each frequency component are less noisy than the original signal. So, when compared to the Fourier-based coherence analysis, the KF-based method should in principle provide sharper mode number estimations.

Experimental results

As a proof of principle, this method was applied to ISTTOK pulse #17081. This is a 3.5 kA pulse (circular plasma) with an edge safety factor $q_a \sim 9 - 10$, line averaged density of $2 \times 10^{18} \text{ m}^{-3}$ and central electron temperature $\sim 100 \text{ eV}$. Magnetic sensor data is acquired at a rate of 2 Msample/s, adequate to cover most of the MHD and turbulence activity, typically in the range $\sim 50 - 300 \text{ kHz}$. The acquired signal window for one of the magnetic sensors used to perform the coherence analysis is shown in figure 3.13. Less than 500 μs worth of data corresponding to less than 1000 samples. Elementary visual analysis indicates dominant component at $\sim 160 \text{ kHz}$.

The KFHE was configured to estimate the in-phases and quadratures of an equally spaced set of components ($\delta f = 2 \text{ kHz}$) in the range $0 - 250 \text{ kHz}$. A Fourier-based coherence analysis using the Fast Fourier Transform (FFT) algorithm, with the same frequency resolution, would itself require 1000 samples thus making it impossible to perform the required windowed averaging. The KFHE was further configured with $\lambda = 10$ and applied synchronously to the signals acquired by two different magnetic sensors located in the same toroidal position and different poloidal position. A sliding window of 64 samples (in this case a 63 sample overlap was used) was used to perform the averaging required by (3.28) in the calculation of the cross and auto-spectrums. Having these, $C_{12}(\omega)$ can be calculated thus enabling also the calculation of $\gamma_{12}(\omega)$, $\psi_{12}(\omega)$ and $m(\omega)$.

The coherency $\gamma_{12}(\omega)$ obtained with this method is shown in figure 3.14(a). The frequency at which the mode is observed (ω for which $\gamma_{12}(\omega) \sim 1$) is indeed in the expected frequency range. The poloidal mode number m is calculated via (3.27) and plotted in the regions where $\gamma_{12}(\omega) > 0.8$, see figure 3.14(b), suggesting the presence of an $m = 2$ rotating mode.

In order to assess the accuracy of the previous estimation, the same coherence analysis was performed using a standard Fourier-based approach and the results are depicted in

⁶Since determining the ensemble average is not feasible when dealing with unique data series, and assuming that the signals are stationary in a certain interval and that the process is reasonably ergodic, the ensemble average can be replaced by a suitable time average.

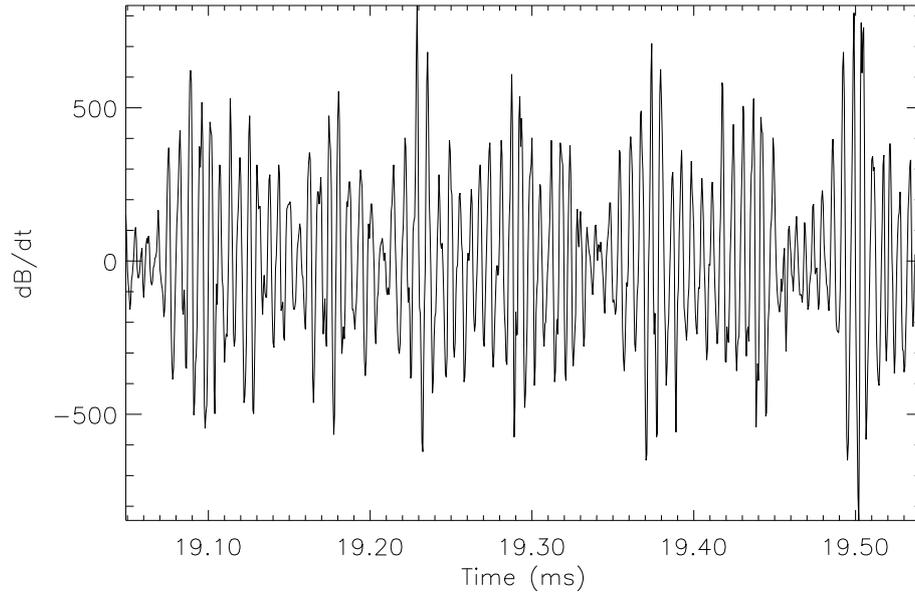


Figure 3.13: Mirnov coil signal ($\theta = 60^\circ$, pulse #17081) evidencing an $m = 2$ mode rotating at ~ 160 kHz

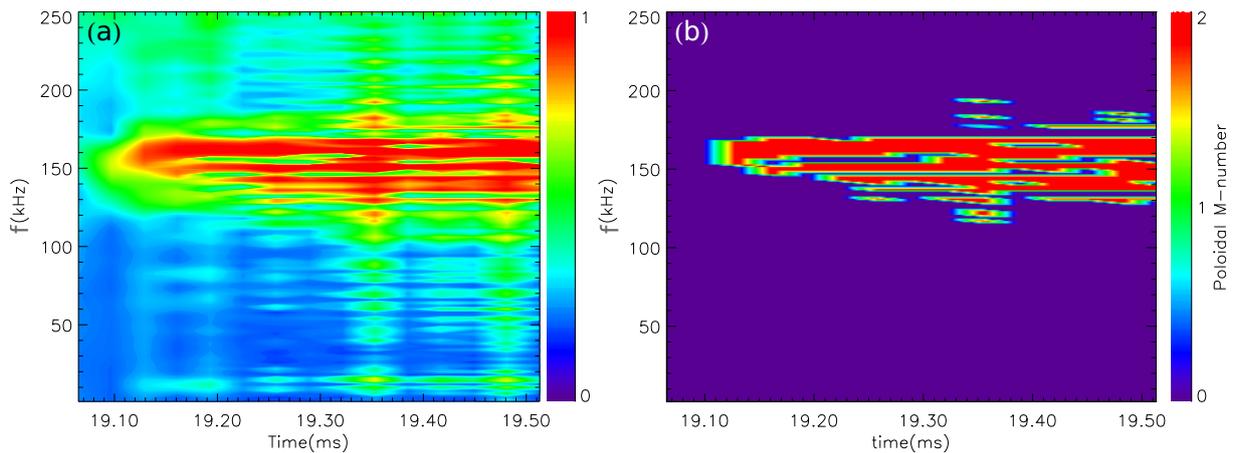


Figure 3.14: KFHE-based - (a) Coherency spectrum and (b) Poloidal mode number spectrum for $\gamma > 0.8$.

figures 3.15(a) and 3.15(b). Cross-power spectral density is performed on sliding macro windows of 196 samples (20% of the total signal length), each separated by 10 samples. For each of these macro windows, 20 sliding DFTs are performed on 116 signal samples (zero-padded with 884 samples thus providing an *apparent* resolution of $F_s/(116+884) = 2$ kHz) separated by 4 samples. Each result is therefore produced every 10 samples ($5 \mu\text{s}$) with a latency of $196/2 = 98$ samples (equivalent to $49 \mu\text{s}$) though the temporal localization is dispersed over the macro window. The KF-based method however, as previously mentioned, uses an equivalent macro window of 64 samples but requires no sub-windows. Both methods identify coherent $m = 2$ magnetic activity in the vicinity of 160 kHz however, it is evident that the limited number of samples (insufficient statistics) has a greater impact on the capabilities of the Fourier-based method to decorrelate incoherent fluctuations.

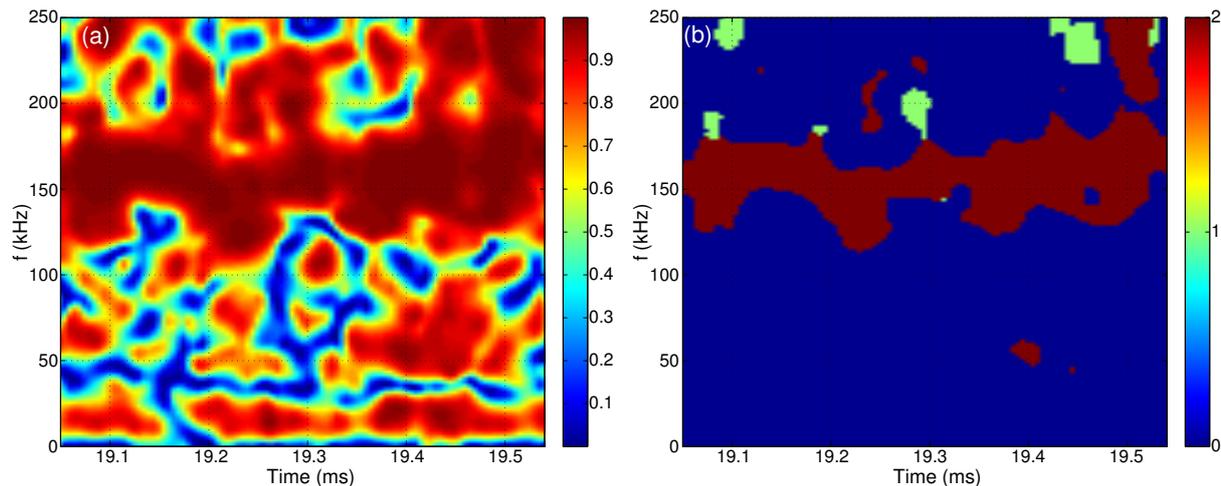


Figure 3.15: Fourier-based - (a) Coherency spectrum and (b) Poloidal mode number spectrum for $\gamma > 0.95$.

3.3.3 TAE signal processing I

Routine studies are performed on JET using a new set of antennas to excite Toroidal Alfvén Eigenmodes (TAE). A TAE resonance footprint is observed in the plasma response measurement when there is a noticeable variation in both the amplitude and the phase of the response with respect to the excitation. An algorithm for real-time identification of TAE resonances, based on a hardware lock-in amplifier, is presently used at the JET tokamak for detecting such variations. The problem of estimating the I - Q characteristics from a known non-stationary frequency mode, with a resonant-like phase response, embedded in a digital signal is addressed herein. A non-stationary linear model is used in a recursive filter implementation of a lock-in amplifier. One proposes it as a viable alternative to hardware synchronous detectors such as the one in use at the JET and compare its performance with standard digital lock-in techniques in terms of bandwidth and phase response under high throughput rates requirements

Introduction

The usefulness of MagnetoHydroDynamic (MHD) spectroscopy markers in the characterization of tokamak plasmas is widely recognized since it adds valuable information on the already broad span of measurements from plasma diagnostics. One particular useful marker targets a particular type of plasma instabilities identified as TAEs [Chen 85]. The resonant wave-particle interaction between these modes and the fusion born alpha particles (He ions) may lead both to the destabilization of the modes [Fu 89] and to the stochastization of the alpha particles orbit, with a consequent particle and energy confinement loss and possible damage to the first wall [Faso 95b]. Assessing the damping/growth rates, frequency and wavenumber analysis of TAEs may therefore provide valuable information on bulk plasma stability and fast particle confinement losses. The conventional approach to carry out the former is to drive the modes with dedicated antennas [Faso 95a], [Snip 04], operating at frequencies within the range where TAEs are expected to propagate, according to some theoretical model. When the antenna resonates with a particular TAE eigenmode in the plasma, there is a significant increase in the plasma response that is clearly observed in the measured signals of relevant diagnos-

Figure 3.16: Typical spectrogram of a Mirnov coil signal showing clearly the frequency sweeps and resonant spots

tics (e.g., magnetic Mirnov probes, soft X-ray tomography, microwave reflectometry or beam emission spectroscopy [Wong 99]) by noticeable variations in both the amplitude and the phase shift with respect to the excitation. In order to measure the amplitude and the phase of selected signal components immersed in noise and compare to the excitation signal, lock-in amplifiers or synchronous detectors are beyond any doubt the natural solution. This inherently Fourier based method is widely used as a precision electronics measurement tool that plays a fundamental role not only in experimental physics, but also in modern science and engineering in general [Mand 94, Noll 91, Woma 83]. In this work, a general overview of the possible implementations of synchronous detection systems for the analysis of the transfer function of resonant excitation of TAEs is addressed. Focusing in more detail on the real-time techniques for estimating the amplitude and phase of the plasma response, a KF [Kalm 60] digital signal processing approach is investigated as a promising candidate to provide very high throughput estimates with flexible noise mitigation capabilities, capable of contrasting the more conventional LPF approaches that are native in a lock-in amplifier either digital or analog implementations. Concentrating on the particular case of TAE resonance excitation, two strategies are proposed for estimating the damping rate for modes with negligible driving sources, depending on whether one may dispose or not of the reference waveform driving the antennas. A novel implementation using the extended Kalman filter is shown where the amplitude of the plasma response is still recovered although no information is known about the exciting signal.

TAE Antenna Diagnostic

The JET tokamak was among the groundbreakers in actively exploring the TAE mode characteristics in hot relevant plasmas [Faso 95a]. Early mode excitation employed the existing set of installed saddle coils that were more routinely used for error field correction. Due to the coil set up configuration, four coils separated toroidally by 90° in both lower

and upper part of the torus (75 m² in total), little flexibility could be achieved in the range of toroidal wavenumbers (n) that could be driven, i.e., was limited to $|n| \leq 2$ depending on the phasing (positive/negative) of each coil. However, probing the full range of foreseen Alfvén eigenmodes (AE) for JET was possible owing to a broadband power amplifier of 3 kW with a frequency span between 30 kHz up to 500 kHz. The peak current and voltage induced in the saddle coils were, respectively, 30 A and 500 V. The driven magnetic field in the plasma core didn't perturb the plasma significantly owing to a low normalized $\delta B/B < 10^{-5}$ where δB represents the magnitude normalized of the induced magnetic field perturbation and B a typical value for the toroidal field $\sim 1 - 3$ T. After 2005, an upgrade to the diagnostic was carried out [Test 04] to allow for a much more compact setup and consequently higher toroidal mode probing. Two sets of four rectangular coils toroidally apart were installed, each made of 18 loops with 25 cm \times 25 cm covering a total area of 0.5 m² in total and placed 4 cm behind the poloidal limiter. This results in a wider range of toroidal modes that can be excited ($|n| < 50$). An upgrade to the broadband amplifier to 5 kW was made and current and voltage peak values of 15 A and 700 V are achieved, yielding a maximum magnetic perturbation in the range $\delta B \sim 10^{-9} - 10^{-8}$ T. Routine studies are performed on JET using this new set of TAE antennas to excite modes typically in the 100-400 kHz frequency range and measure the plasma response with Mirnov coil signals [Test 07, Klei 08b]. This identification mechanism is part of a closed feedback loop that controls the antennas' sweeping frequency direction in order to maximize the number of resonant crossings during a JET pulse. The frequency of the driving magnetic perturbations is controlled by a Voltage Controlled Oscillator (VCO) with characteristic sweeping values of 200 kHz/s setting the time and frequency resolution of the measurements for the damping rate and resonant frequency. This is shown in figure 3.16 for JET pulse #69571 where the resonance matching is clearly evidenced in the spectrogram from a Mirnov signal. The spectrogram, however, provides only a first insight of where the resonance lies in frequency space but not on the actual damping rate of the driven mode. A practical estimate, adequate for a real-time implementation, for the damping rate may be derived from the Full Width at Half Maximum (FWHM) of the measured amplitude response as a function of the probing frequency and the resonant frequency is simply obtained from the maximum response. Alternatively, more refined measurements are possible from the transfer function $H(\omega) = M(\omega)/R(\omega)$, where M and R stand, respectively, for the complex valued plasma response and the reference excitation. Resonances manifest themselves as poles in the transfer function, from which the damping coefficient can be extracted [Faso 95a].

Synchronous detection basically attempts to identify the components of the plasma response, with the same frequency as the excitation, that are in-phase and in quadrature (I - Q) with the excitation. In a complex plane representation, I and Q represent the real and imaginary components of the response. Far away from the resonance, I and Q should remain constant thus indicating no change in both amplitude and phase while crossing the resonance eventually leads to a circular pattern in the complex plane.

At JET, the real-time identification of TAE resonances is based on a hardware INCAA based synchronous detection system with up to 48 channel boards. The I - Q components are obtained by analog mixing of the measured signal with in-phase and quadrature references and applying an analog LPF with ~ 100 Hz bandwidth. Although a hardware implementation for the synchronous detection is relatively straightforward, a software approach, in principle more flexible, less expensive and requiring less maintenance, is less trivially implemented.

Standard Methods for Synchronous Detection

As illustrated in the previous section, the quality of the TAE antenna diagnostic measurement is critically dependent on the effectiveness of the synchronous detection block. The goal is to estimate the amplitude and phase response of a system that exhibits a resonant behavior when subject to a known external excitation under certain conditions. The current hardware implementation of the synchronous detection can be equivalently done in software by means of a Fourier projection of the plasma response (Mirnov pick-up coil signal and quadrature) onto the excitation reference and its corresponding quadrature with normalized unitary amplitude. Incidentally, the synchronous detection procedure just described provides the basic principle of a DPL, see section 3.3.1. Alternatively, a less biased estimate of the amplitude may be obtained by mixing the plasma response with a delayed version of the excitation reference, in-phase with the former and with normalized unitary amplitude. The subsequent low-pass filtering of this mix yields the final amplitude estimation in what is known as a SPL, see also section 3.3.1.

This homodyne-type detection requires a large enough amount of samples to be used for the LPF to effectively filter out noise and provide the in-phase and quadrature components. These synchronous detectors are powerful methods although one may note some inherent shortcomings: noise mitigation may require large (order 100) of samples in an FIR implementation of an LPF, thereby increasing the *processing burden* and inevitably introducing a time delay in the amplitude and phase estimates; it is best suited when there are no restrictions on the amount of samples used for the LPF FIR; real-time tuning of the filter parameters for balancing the width of the low-pass region and the delay, although possible, is inherently an intricate problem since it requires changing in real-time, not only the value of the filter coefficients, but also their total number. Alternatively, one may opt for IIR filters. The performance of both types of filters is similar in terms of magnitude response and introduced delay. Although IIR filters require a great deal less coefficients than the FIR filters for similar magnitude and phase responses, a particular attractive feature for real-time implementations, this comes at the price of, unlike the FIR filters, having a non-linear phase response. This, in practice, means that the IIR filters introduce different delays depending on the frequency components present in the signal and lead to distortions in the estimated signals.

The Kalman filter synchronous detector

For the purpose of this work, the KFHE is reused in a single-tone implementation to characterise the plasma response to the magnetic excitation. The novelty here, compared for instance to what was done in the case of the MSE signal processing, is that the frequency component that is required to be tracked is non-stationary but its evolution in time is predefined and feed-forwarded to the VCO in real-time. Therefore, one has the knowledge, at all times, of the frequency request to the VCO but may or may not have the actual reference (excitation) signal fed to the antennas available. As previously mentioned, resonance crossings can be identified by noticeable and characteristic changes in both the amplitude and relative phase difference in the Mirnov signal. In particular, one of the quantities of interest is the excited modes' damping rate. For this purpose one can either use the phase response, the amplitude response or both in a complementary way. Extracting the instantaneous amplitude of the signal component of interest is relatively straightforward (the amplitude at time sample k is $a_k = \sqrt{x_{1,k}^2 + x_{2,k}^2}$). It does not require the reference signal to be available for processing but it does require real-time

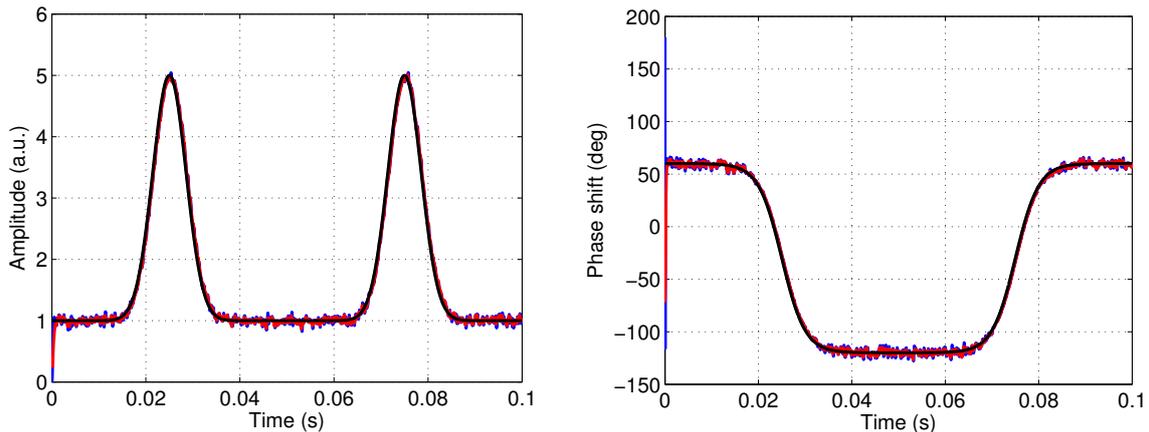


Figure 3.17: The DPL estimate (blue), KFHE estimate (red) and true value (black) of: (a) amplitude of the plasma response and (b) phase difference between excitation and plasma response.

knowledge of its frequency evolution. By analyzing the FWHM of the resonant peak it is also possible to estimate the damping rate. If however one requires an estimate of the phase response as a cross check, then it is critical to also have the in-phase and the quadrature components of the reference signal, i.e., $x_{1,k}^{ref}$ and $x_{2,k}^{ref}$ respectively. Provided these are available, one can immediately build the in-phase projection $I \propto \cos(\Delta\phi(\omega))$, quadrature projection $Q \propto \sin(\Delta\phi(\omega))$ and extract the phase difference $\Delta\phi(\omega)$ between the reference excitation and the plasma response at the given frequency ω . In case both the reference signal and the frequency requested to the VCO are available in real-time, one can estimate the quadrature component by using the single-tone KFHE directly on the reference signal. This is a good advantage of the Kalman filter approach since it contributes to the overall simplification of the estimation process.

Unfortunately, with the present setup of the TAE antenna diagnostic at JET the VCO output signal, which is essentially the excitation waveform that drives the antennas, is not stored for post-pulse analysis. Therefore, in order to perform a fair comparison between the DPL and the KFHE, the excitation and the plasma response signals have been synthesized with an amplitude and phase evolution as displayed, respectively, in figure 3.17(a) and figure 3.17(b) as indicated by the black line. The input frequency of the excitation is swept with a trend similar to the one used in the real experiments as shown in figure 3.18. The sampling frequency was set to 1 Msample/s and normally distributed random noise yielding 0.5 signal-to-noise ratio was added to the synthetic Mirnov signal. As shown in figure 3.17(a) and figure 3.17(b), the performance obtained with the KFHE for $\lambda = 2 \times 10^4$ is comparable with the one obtained with the DPL with a FIR LPF with 3 kHz cutoff frequency and 400 taps (200 sample average delay). Amplitude and phase estimations are provided in this case at a rate of 20 ksamples/s. It is worth mentioning that, although a IIR LPF would require less filter coefficients, its non-linear phase response may introduce non-trivial and non-acceptable signal distortion.

In many real-time systems running at a predefined cycle timing interval ($1/F_{out}$) a single estimation is often required to be derived entirely from the N samples acquired within the cycle itself ($N = F_s/F_{out}$) where F_s and F_{out} are, respectively, the sampling frequency and the real-time network throughput rate). For the particular case where a 50 kHz cycle time is required implying that only 20 samples are available each cycle, figure 3.19

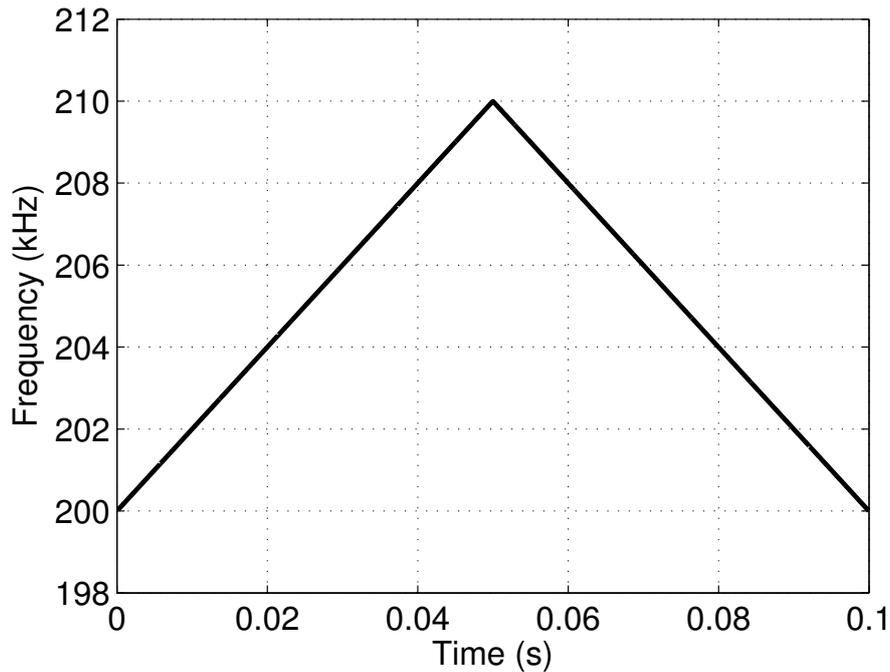


Figure 3.18: Frequency evolution of the excitation waveform.

shows the comparison between the amplitude estimations provided by the KFHE and a DPL with a 20 tap (10 sample delay) FIR LPF. The DPL obviously has a lower delay than the KFHE at the expense of noisier estimates. Also, whereas in the case of the DPL a change in the bandwidth/delay requires recalculation of the filter coefficients making it hard for performing multi-rate signal processing in real-time, in the case of the KFHE the only requirement is to change the λ factor since the filter depends only on the estimate of the previous system state.

The KFHE provides relevant improvements over the DPL mainly in the cases when a predefined and fast cycle time imposes the usage of only a few samples for the LPF although its robustness and ease of use make it essentially a more maintainable tool.

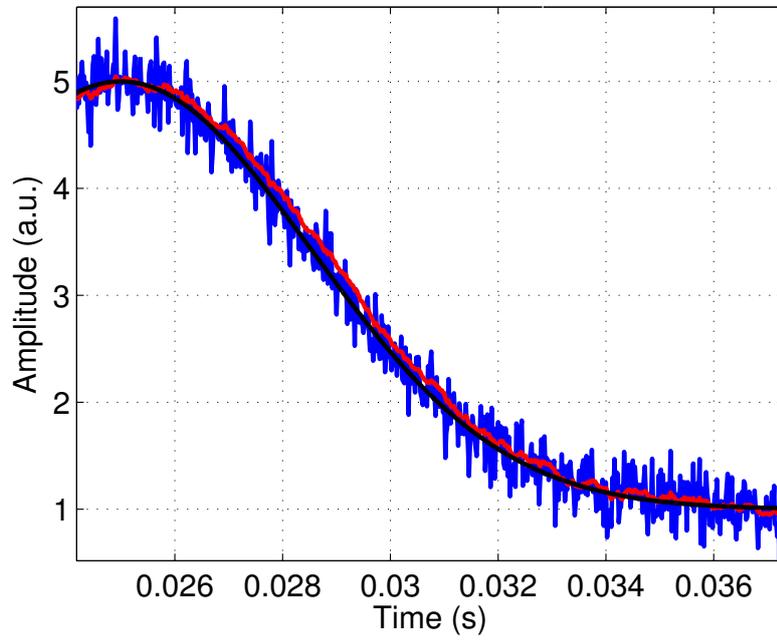


Figure 3.19: Amplitude estimation - DPL estimate with 20 samples for LPF-ing (blue), KFHE estimate (red) and true amplitude (black).

Chapter 4

Non-stationary and non-linear frequency tracking and harmonic estimation

4.1 Introduction

The non-linear KF in the FT configuration, proposed by La Scala and Bitmead in an EKF implementation [La S 96], was first applied to nuclear fusion research as a proof of principle in plasma resonance detection for the TAE antenna diagnostic [Alve 10] at JET. A more systematic and detailed study of the potential of this approach was done by Alves and Coelho [Alve 11a] demonstrating its competence as a real-time algorithm for the simultaneous estimation of the instantaneous amplitude, frequency and phase of multiple non-stationary signal components. This following section will be devoted to the description of the FT configuration that will be implemented in various non-linear KF flavours.

Developments presented herein are based on [Alve 10, Alve 11a, Alve 13a, Alve 13b] for which the author of the present dissertation provided a substantial contribution.

4.2 Theory

Assuming a generic signal S_k , see equation (3.1), where k denotes the time index, to be the sum of N non-stationary signal components $S_{i,k}$ exhibiting the properties of an intrinsic mode function [Huan 98], i.e., $S_{i,k} = A_{i,k} \sin(\phi_{i,k})$ (where $A_{i,k}$ is the instantaneous amplitude and $\phi_{i,k}$ is the instantaneous phase) plus a noise term n_k , then \mathbf{f}_k describes the non-linear frequency tracking model/transformation for a single signal component, see equation (4.1). Rewriting the previous in state space notation yields equation (4.2), where $\hat{\mathbf{x}}$ is the prior state estimate and \mathbf{x} is the posterior state estimate, illustrating that the model is a rotation in the I-Q plane with phase increment $x_{3,k-1}$.

$$\mathbf{f}_k : \mathbb{R}^3 \rightarrow \mathbb{R}^3 : (x_{1,k}, x_{2,k}, x_{3,k}) \rightarrow \begin{pmatrix} x_{1,k-1} \cos(x_{3,k-1}) - x_{2,k-1} \sin(x_{3,k-1}), \\ x_{1,k-1} \sin(x_{3,k-1}) + x_{2,k-1} \cos(x_{3,k-1}), \\ x_{3,k-1} \end{pmatrix}, \quad (4.1)$$

$$\begin{bmatrix} \widehat{x}_{1,k} \\ \widehat{x}_{2,k} \\ \widehat{x}_{3,k} \end{bmatrix} = \begin{bmatrix} \cos(x_{3,k-1}) & -\sin(x_{3,k-1}) & 0 \\ \sin(x_{3,k-1}) & \cos(x_{3,k-1}) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{1,k-1} \\ x_{2,k-1} \\ x_{3,k-1} \end{bmatrix} \quad (4.2)$$

For a single tone, $x_{1,k}$ represents the signal estimate, $x_{2,k}$ can be viewed as the real-time Hilbert transform of $x_{1,k}$ (i.e. the quadrature of $x_{1,k}$) and $x_{3,k}$ is the signal's instantaneous phase increment. In some senses $x_{1,k}$ can be interpreted as the result of a non-stationary band-pass filter around the tracking signal component. Furthermore, if the only observable quantity is a direct measurement of the signal S_k , then the measurement equation is linear and, for a single tone, given by (4.3).

$$\widehat{z}_k = \mathbf{C}_k \widehat{\mathbf{x}}_k = [1 \ 0 \ 0] [\widehat{x}_{1,k} \ \widehat{x}_{2,k} \ \widehat{x}_{3,k}]^T \quad (4.3)$$

The real-time estimations of instantaneous frequency, amplitude and phase, for a single non-stationary signal component are given, respectively, by equations (4.4), (4.5) and (4.6) where F_s is the measurement's sampling frequency.

$$f_k = \frac{x_{3,k}}{2\pi} F_s \quad (4.4)$$

$$a_k = \sqrt{x_{1,k}^2 + x_{2,k}^2} \quad (4.5)$$

$$\phi_k = \arctan \left[\frac{x_{2,k}}{x_{1,k}} \right] \quad (4.6)$$

Generalising the FT equations for the case of multi-component analysis is straightforward and requires block diagonally stretching the transformation matrix of equation (4.2), block stretching the measurement matrix of equation (4.3) and appropriately adjusting the dimensions of all the other matrices. Therefore, the straightforward adaptation of (4.1) and (4.2) to the multi-component case yields a new process matrix and a new measurement equation given respectively by (4.7), where $A'_{i,k}$ is given by (4.8), and (4.9).

$$A_{N_h,k} = \begin{bmatrix} A'_{1,k} & & & \\ & A'_{2,k} & & \\ & & \ddots & \\ & & & A'_{N_h,k} \end{bmatrix} \quad (4.7)$$

$$A'_{i,k} = \begin{bmatrix} \cos(x_{3i,k-1}) & -\sin(x_{3i,k-1}) & 0 \\ \sin(x_{3i,k-1}) & \cos(x_{3i,k-1}) & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad i = 1 \dots N_h \quad (4.8)$$

$$\begin{aligned} \widehat{z}_k &= \mathbf{C}_k \widehat{\mathbf{x}}_k \\ &= [1 \ 0 \ 0 \ 1 \ 0 \ 0 \ \dots \ 1 \ 0 \ 0] [\widehat{x}_{1,k} \ \widehat{x}_{2,k} \ \widehat{x}_{3,k} \ \widehat{x}_{4,k} \ \widehat{x}_{5,k} \ \widehat{x}_{6,k} \ \dots \ \widehat{x}_{3N_h-2,k} \ \widehat{x}_{3N_h-1,k} \ \widehat{x}_{3N_h,k}]^T \end{aligned} \quad (4.9)$$

The generalized real-time estimates of the instantaneous amplitude and phase, for each component, is given respectively by (4.11) and (4.12).

$$x_{3i,k} = 2\pi f_{i,k} / F_s \quad (4.10)$$

$$a_{i,k} = \sqrt{x_{3i-2,k}^2 + x_{3i-1,k}^2} \quad (4.11)$$

$$\psi_{i,k} = \tan^{-1} \left(\frac{x_{3i-1,k}}{x_{3i-2,k}} \right) \quad (4.12)$$

Finally, under the reasonable assumptions that $\mathbf{R} = r\mathbf{I}$ and $\mathbf{Q} = q\mathbf{I}$ where r and q are scalars and \mathbf{I} is the identity matrix, the asymptotic value of the KF gain K is a function of the ratio $\lambda = r/q$ rather than the values of r and q independently [Bitt 00].

4.3 Studies

In order to assess the behavior of the EKFFT as proposed in [La S 96], a series of tests have been performed under a variety of different circumstances. For all these tests one assumes that the original input signal is given by equation (3.1), where N_h is the total number of tones and n_k is a time series of normally distributed random noise. By *tone* one means any, stationary or non-stationary, signal component with the properties of an intrinsic mode function [Huan 98]. The instantaneous quantities: frequency, amplitude and phase are given, respectively, by equations (4.4), (4.5) and (4.6).

Q is a $3N_h \times 3N_h$ diagonal matrix which we have set to be $Q = Q_{diag} = [q_1^a \ q_1^a \ q^f \ q_2^a \ q_2^a \ q^f \ \dots \ q_{N_h}^a \ q_{N_h}^a \ q^f]$. As noted in [Coel 09a], when using a linear harmonic estimator with a stationary frequency grid and uncorrelated process noise ($Q = qI$ with I the identity matrix) the filter response depends on the ratio R/q alone. In the case of non-stationary harmonic estimation, and as described in [Bitt 00], the same behavior applies. However, some care should be taken when heading towards different process covariances for the signal (q_i^a) and frequency (q^f) estimations. Since the bandwidth of the filter response is actually tied to the covariance of the signal estimation and much less to the covariance of the frequency estimation, one should naturally expect that, if the frequency estimation at some time instant is very far from the true value, the outcome in terms of signal estimation will also be very poor. Therefore, it is widely anticipated that the covariance one sets for the frequency estimation (q^f) should be much smaller than the signal counterpart (q_i^a). This will enforce the estimated frequency to have small excursions in time while relaxing the signal estimation excursions, thereby improving the adherence to the underlying signal component.

4.3.1 Single Tone

Starting with the most basic of all tests, and in order to perform a preliminary assessment of the EKFFT's capabilities, one has synthesized a single tone that comprises a mix of AM (10 ± 5 a.u. @ 12.5 Hz) and FM (15 ± 2.5 kHz @ 50 Hz) with additive zero mean Gaussian distributed random noise ($\sigma = 0.25$). The sampling frequency used is $F_s = 1$ MHz. It can be seen from figure 4.1 that the EKFFT is able to track, in real-time, both the instantaneous amplitude and frequency evolution, of the single tone embedded in a noisy signal, with a phase distortion inferior to 10 degrees. This achievement was attained despite the initial frequency guess of 50 kHz and using $q_1^a = R = 1$ and $q^f = 10^{-5}$.

Because equation (4.13) is true for any tone (intrinsic mode function [Huan 98]), i , not only (4.4) but also (4.14) is eligible to provide estimates of the instantaneous frequency. Of course, both approaches can only provide the same exact estimates if either their

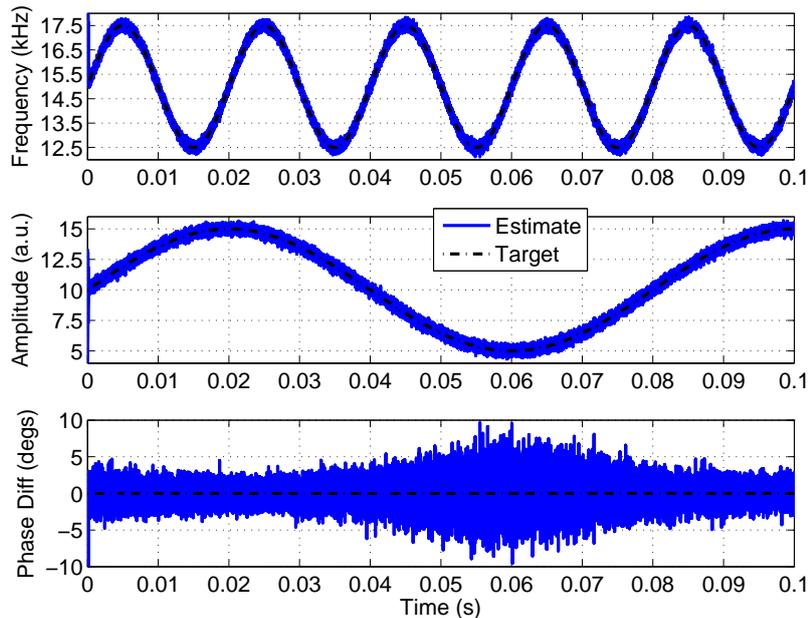


Figure 4.1: Estimation of a non-stationary single-tone in a noisy signal

bandwidth is the same¹ or the pass-band of the estimation using (4.14) is a Dirac- δ function centred in f_k . In practice, when tuning the EKFFT, the aim is to have (4.4) as a smoothed version of (4.14).

$$\omega_i(t) = \frac{\partial \phi_i(t)}{\partial t} \quad (4.13)$$

$$f_k = \frac{1}{2\pi} \frac{\partial}{\partial t} \left[\tan^{-1} \left(\frac{x_{2,k}}{x_{1,k}} \right) \right] \quad (4.14)$$

4.3.2 Amplitude Separation

In this section one presents the first test of a simultaneous two tone estimation with the EKFFT. The aim is to establish, for two stationary sinewaves separated by 5 kHz, the estimation performance of the EKFFT as a function of the amplitude ratio of the signal components. Each tone, i , is given by equation (4.15), where $a_1 = 1$, $f_1 = 15$ kHz and $f_2 = 20$ kHz. The total signal is given by $S(t) = S_1(t) + S_2(t) + n(t)$, where $n(t)$ is normally distributed random noise with $\sigma = 0.05$.

$$S_i(t) = a_i \times \sin(2\pi f_i t) \quad (4.15)$$

For this first test, one has setup the EKFFT in the multi-tone configuration (simultaneous estimation of the $f_1 = 15$ kHz and $f_2 = 20$ kHz frequency tones), provided the correct initial frequency guesses and let $q_1^a = q_2^a = R = 1$ and $q^f = 10^{-5}$. Figure 4.2 shows the resultant time traces of the EKFFT's amplitude estimation, of the $f_1 = 15$ kHz component, for different values of a_2 . It can be seen that for $a_2/a_1 \geq 100$ the EKFFT performs poorly because $S_2(t)$ is such a strong component that it masks out the presence

¹difficult to achieve due to the non-linearity of the process although, in principle, possible by adjusting q_i^a and q^f .

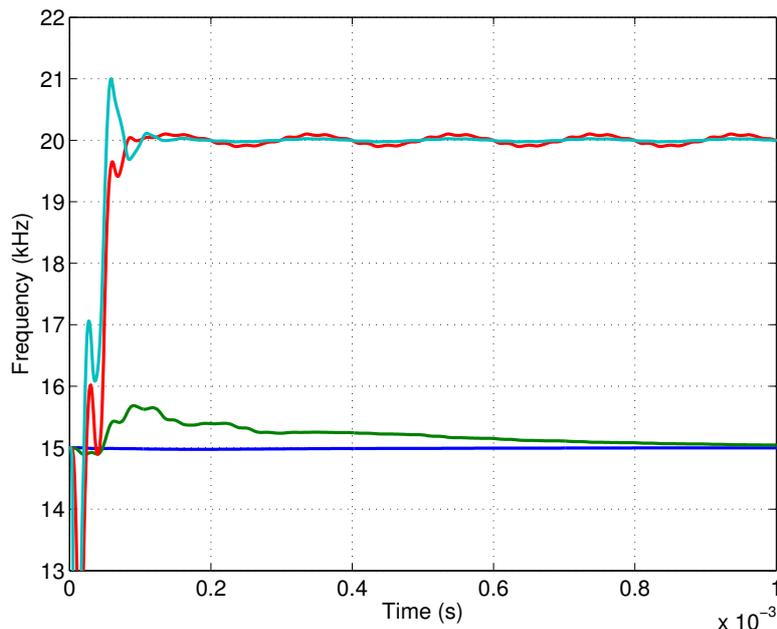


Figure 4.2: Estimation of the frequency of the 15 kHz component for the cases where a_2 is 10^0 (blue), 10^1 (green), 10^2 (red) and 5×10^2 (cyan).

of $S_1(t)$. It is clear that, in order to improve the performance of the EKFFT in this sort of cases, one requires a way of tuning the EKFFT by changing its rate of belief in either the measurement or the model as a function of each component's amplitude. In order to do so, one starts by recalling that, as shown in [Bitt 00], the performance of the EKFFT depends only on the ratio $\lambda = R/q$ where $\mathbf{Q} = q\mathbf{I}$ and \mathbf{I} is the $3N_h \times 3N_h$ identity matrix. For all tests one has used $\mathbf{Q} = \mathbf{Q}_{diag} = [q_1^a \ q_1^a \ q^f \ q_2^a \ q_2^a \ q^f \ \dots \ q_{N_h}^a \ q_{N_h}^a \ q^f]$ with $q^f = 10^{-5}$. This choice has proven, through usage experience, to provide robust results regarding diverging problems of the EKFFT for a wide range of ratios $\lambda = R/q_i^a$. Essentially, whilst the ratio R/q^f sets the bandwidth estimation of $x_{3,k}$, the ratio R/q_i^a defines the bandwidth of the band-pass filter centred at f_k given by (4.4). With q^f fixed, the EKFFT's performance depends solely on the ratio R/q_i^a . Considering (3.1) and uncorrelated sinusoidal signal components $S_i(t)$ of the form (4.15), the expression (4.16) for the variance of $S(t)$ holds.

$$\text{var}[S(t)] = \frac{1}{2} \sum_{i=1}^N a_i^2 + \text{var}[n(t)] \quad (4.16)$$

Given the significance of R and Q , and acknowledging that from the point of view of estimating component i all other signal components are at the same level as noise in the sense that they can be viewed as contamination sources, one can intuitively write equation (4.17), where χ represents a lower bound for the ratio q_i^a/R . One easily verifies that, in the limit when $a_i \rightarrow 0$, then $q_i^a/R \rightarrow \chi$ and in the limit when $a_i \rightarrow +\infty$, then $q_i^a/R \rightarrow +\infty$ hence instructing the EKFFT to trust the measurements rather than

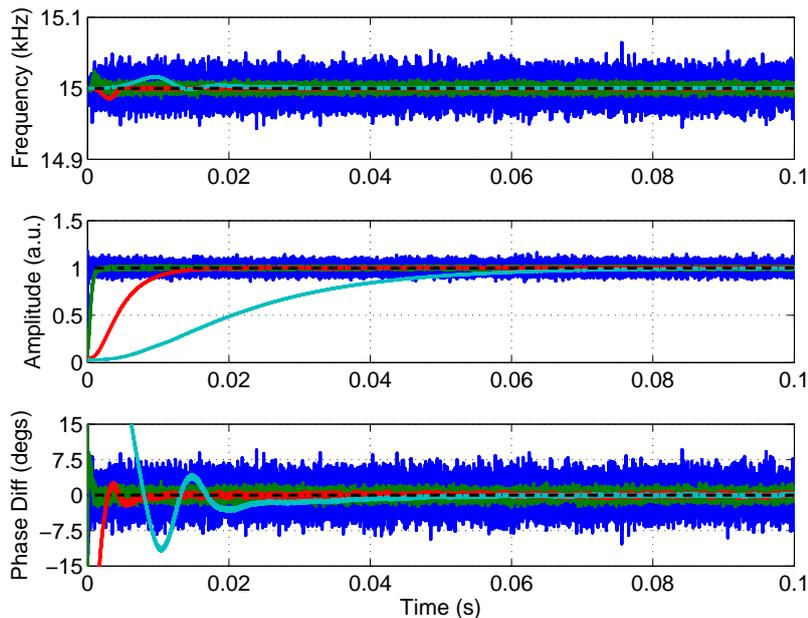


Figure 4.3: Estimates of the 15 kHz signal component in the cases where a_2 is 10^0 (blue), 10^1 (green), 10^2 (red) and 5×10^2 (cyan).

the model for this particular tone.

$$\frac{q_i^a}{R} = \frac{\text{var}[S(t)]}{\frac{1}{2} \sum_{j=1, j \neq i}^N a_j^2 + \text{var}[n(t)]} \chi \quad (4.17)$$

In cases where the variance of the noise is much smaller than the variance of any signal component being estimated (4.17) simplifies to (4.18), which is rather convenient since in many cases one does not know, in real-time, the variance of the noise.

$$\frac{q_i^a}{R} \approx \frac{\sum_{i=1}^N a_i^2}{\sum_{j=1, j \neq i}^N a_j^2} \chi \quad (4.18)$$

When repeating the previous tests using (4.18), immediate improvements can be observed. In figure 4.3 the estimation results for $S_1(t)$ can be seen. As the amplitude ratio a_2/a_1 increases the q_1^a/R ratio decreases and the EKFFT gradually trusts the theoretical model for $S_1(t)$ better than the actual measurements hence, the signal-to-noise ratio increases but the filter's sensitivity to transients (response time) decreases. Globally, after the inevitable convergence time, $S_1(t)$ is correctly estimated.

In figure 4.4 one can see the estimation results for $S_2(t)$. As the amplitude ratio a_2/a_1 increases the q_2^a/R ratio also increases and the EKFFT gradually trusts the measurements better than the theoretical model for $S_2(t)$ hence, the filter's response time also increases. Similarly, as a_2 increases, $\text{var}[S_2(t)]$ also increases with respect to $\text{var}[n(t)] = 0.05^2$ thus improving the signal-to-noise ratio.

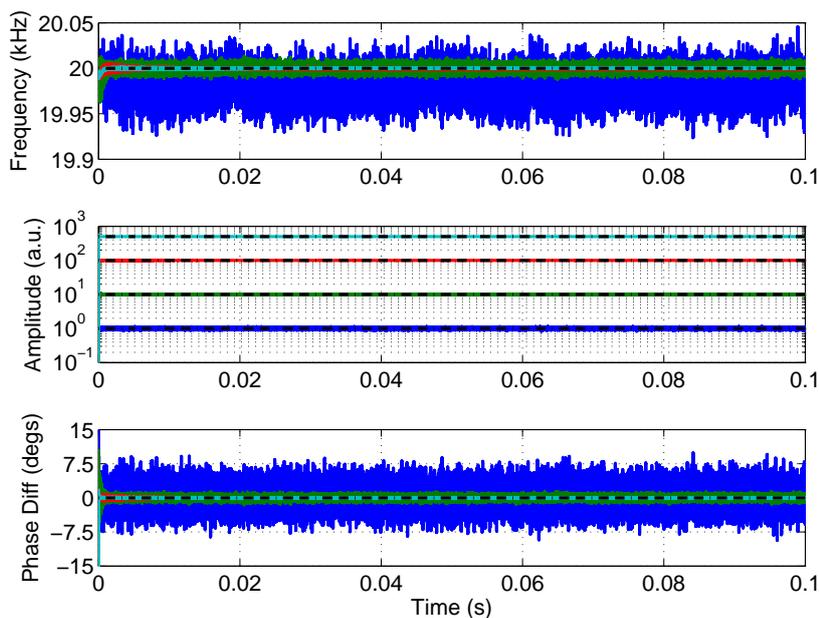


Figure 4.4: Estimates of the 20 kHz signal component in the cases where a_2 is 10^0 (blue), 10^1 (green), 10^2 (red) and 5×10^2 (cyan).

From these results one can conclude that multi-tone estimation in the presence of signal components with amplitude ratios differing by two or even three orders of magnitude is possible depending also on their frequency separation as will be addressed in the following section.

4.3.3 Frequency Separation

In this section the issue of estimating, simultaneously, two tones whose frequency separation is evolving in time is addressed. In particular, one is interested in understanding the behaviour of the system in the limit where the frequencies of two different components become gradually closer until eventually coinciding. For this purpose, two tests were devised. The signal to be analysed is given by equation (4.19), where $n(t)$ is a Gaussian distributed time series representing noise. In the case of the first test, whose results are shown in figure 4.5, one observes that, as the descending frequency chirp gets closer to the stationary 15 kHz component, and although the frequency estimations remain quite good, both the amplitude and phase locking of the components estimates' become biased. In this case, where $a_2 = 1$ and $n(t) = 0$, the phase distortions are always below 20° but clearly there is a performance impoverishment which is worst in the transitions when both frequencies first meet and later when they separate.

$$S(t) = \sin(2\pi f_1 t) + a_2 \times \sin(\phi_2(t)) + n(t) \quad (4.19)$$

Figure 4.6 shows the results for the same exact simulation only with $a_2 = 10$ and $n(t)$ Gaussian noise with $\sigma = 0.05$. We can see that the global behavior of the system is quite similar to the previous test however, the weaker component is notoriously more affected by the presence of the stronger one than in the previous example.

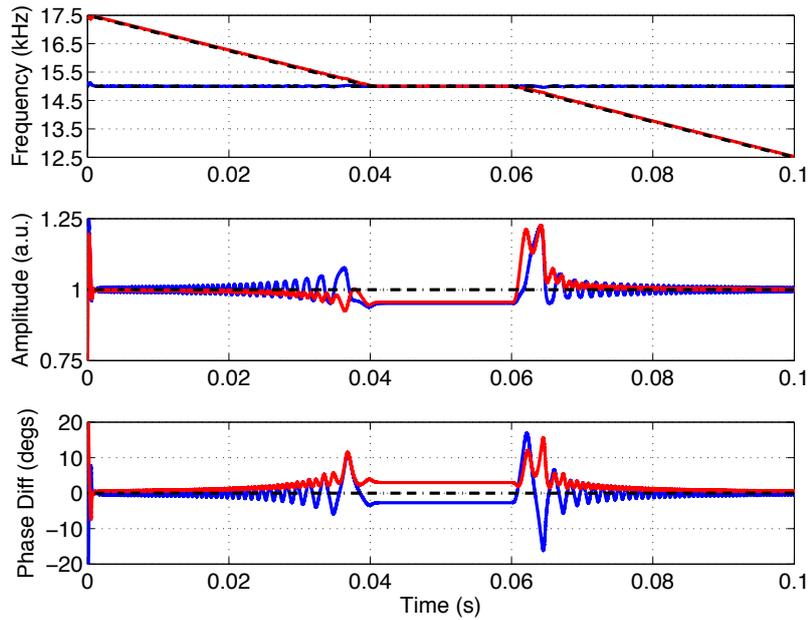


Figure 4.5: Estimates of two signal components of equal amplitudes.

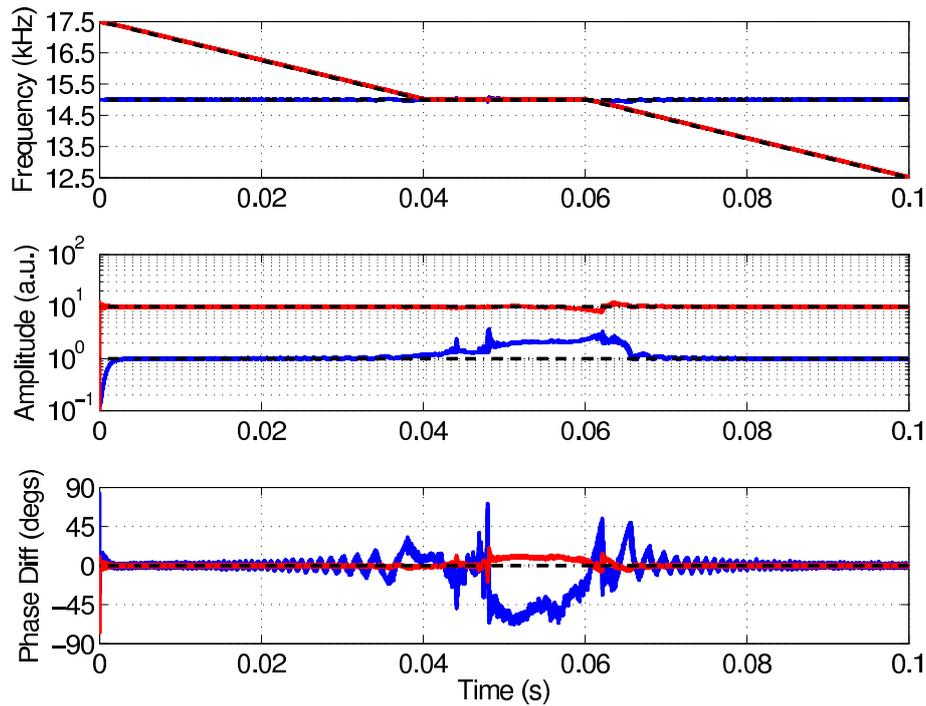


Figure 4.6: Estimates of two signal components of different amplitudes in a noisy signal.

In summary, and as expected, the EKFFT reveals a performance deterioration as the relevant signal component's frequency separation approaches zero. Under these conditions, and as one would intuitively predict, the best performance of the EKFFT occurs when both tones of similar frequency have equal amplitudes. In particular, and as a rule of thumb, we might set as a frequency separation limit, for the case when both tones have equal amplitude, the point where the phase distortion of the signal estimation becomes

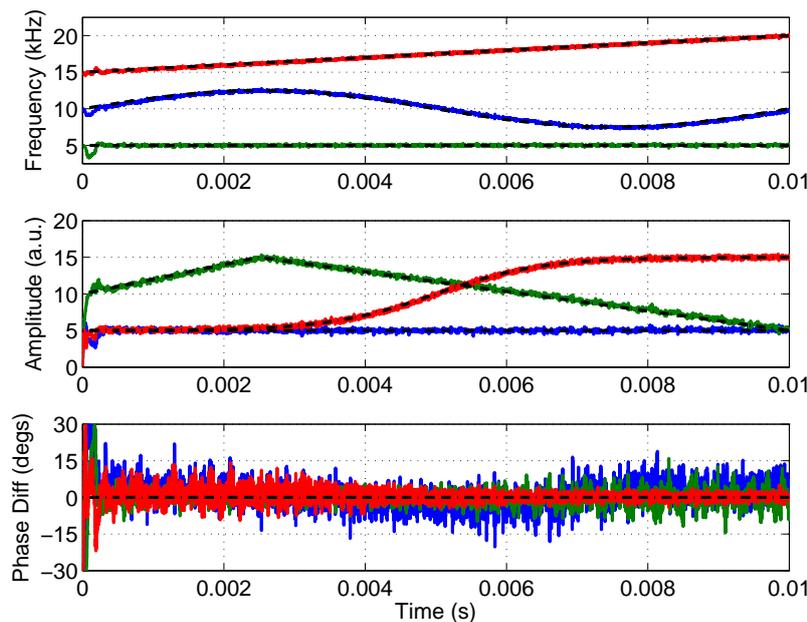


Figure 4.7: Multi-tone estimation in a noisy signal.

larger than 5 degrees. This occurs, for the example depicted in figure 4.5, for a frequency separation of approximately 500 Hz, i.e., 10^{-3} if normalised to the Nyquist frequency.

4.3.4 Multi-Tone Harmonic Estimation and Frequency Tracking

In this section the results of a more complex test are presented. Three different non-stationary signal components are present in a Gaussian distributed noisy ($\sigma = 0.25$) signal, as shown in figure 4.7. In this case the amplitude ratio between components does not exceed a factor of 3 and the lowest frequency separation, normalised to the Nyquist frequency, is 5×10^{-3} . It is interesting to note the increase in the phase difference between the in-phase estimate and the corresponding signal component when its amplitude is reduced (thus is more affected by the signal noise), and the converse when its amplitude increases. Again, the frequency of every component at $t = 0$ is assumed to be known and, in this case, the adaptive correction for the q_i^a/R ratio is not used as the amplitude excursion among signal components is not too demanding. It is clear though, that the EKFFT performs remarkably well in this case where there is the requirement for the simultaneous estimation of three non-stationary (not only in terms of frequency but also amplitude) tones within a signal corrupted with noise.

4.3.5 Remarks

An adaptive method for performing spectral-like real-time analysis of non-stationary components in noisy signals has been proposed. An interpretation of the filter's behaviour, regarding both signal and frequency estimations in terms of bandwidth and its closely coupled relationship with the balance between measurement and model covariances, has also been presented. It is important to stress that in all simulations presented, the estimations' output rate is the same as the measurement's sampling rate which means that,

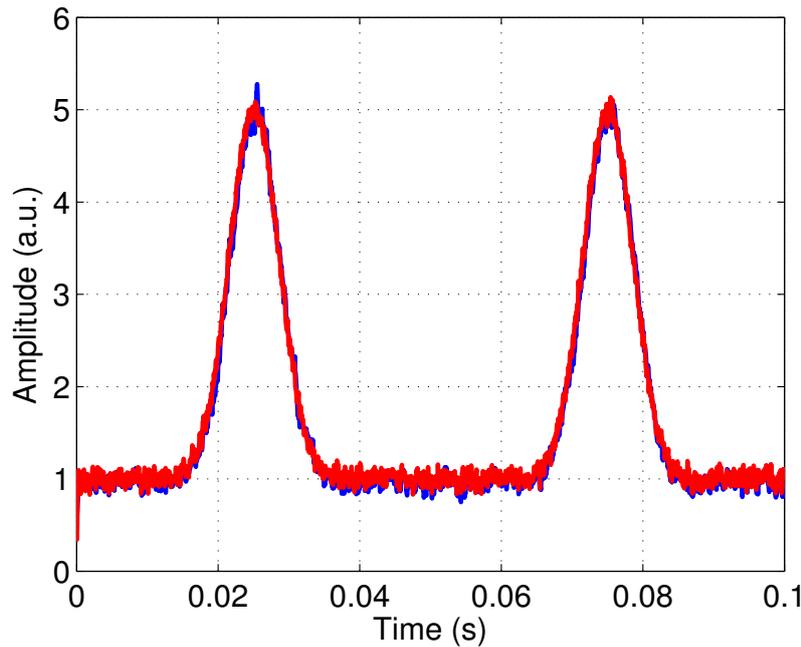


Figure 4.8: Amplitude estimation performance comparison between the EKF (blue) and KF (red).

unlike Fourier-based methods, no filtering/windowing is required although, of course, possible.

Throughout section 4.3, and apart from the test presented in section 4.3.1, at least some basic knowledge of the underlying signal was assumed. In particular, in section 4.3.2, not only one started the estimation process providing the correct frequencies as initial guesses but, for the q_i^a/R adaptation, prior knowledge of the component's amplitudes was also assumed. Regarding the latter, and since the EKFFT already provides the instantaneous amplitudes for each component, it might be possible to envisage a self contained scheme where the real-time adaptation of the q_i^a/R ratio is performed exclusively using information generated within the EKFFT itself. In sections 4.3.2, 4.3.3 and 4.3.4 additional prior knowledge on the amount of relevant signal components was also assumed. It is a known fact that a good quality estimation of a particular component requires the simultaneous estimation of all other signal components carrying relevant amplitudes with respect to the magnitude of the total signal. However, and unless there is a particular interest in a specific tone, estimating tones of minor relative weight with respect to the total signal magnitude will unnecessarily increase the processing burden whilst providing little or no revenue at all in terms of filtering performance.

4.4 Applications

4.4.1 TAE signal processing II

Following the study presented in section 3.3.3, in the case when the amplitude estimation suffices for the calculation of both the resonant frequency (peak in the amplitude response) and the damping rate (given by the FWHM), the reference waveform need not be known because it is only relevant for calculating the systems' phase response with respect to

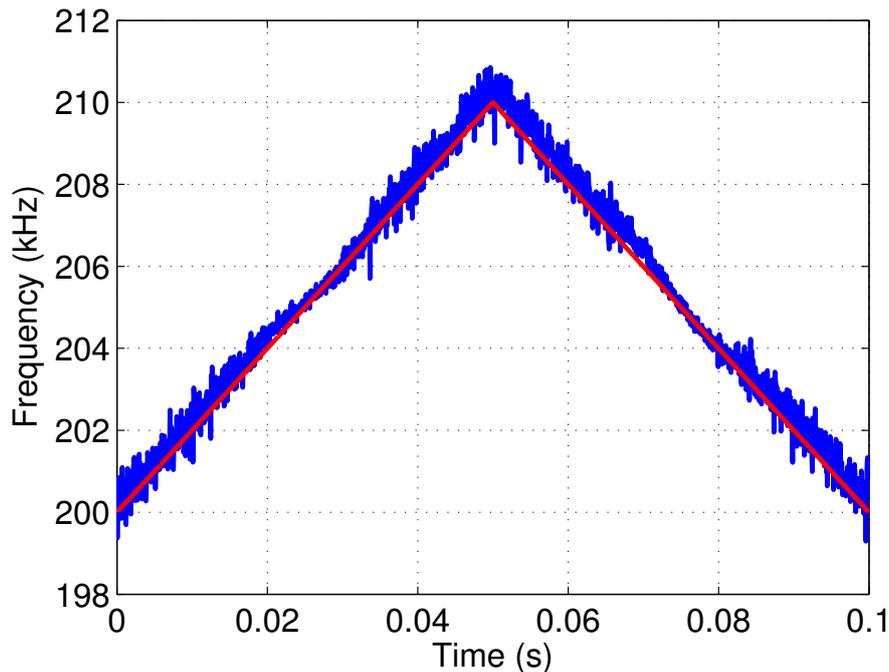


Figure 4.9: EKF frequency estimation (blue) and true frequency content (red).

the excitation. In figure 4.8, the amplitude estimation performance comparison between the KFHE and the EKFFT is shown. While the KFHE requires the online knowledge of the signal frequency but does not require a reference signal, the EKFFT not only needs none of that information but also provides its own estimation of the signal frequency (see figure 4.9) with a similar signal-to-noise ratio and time delay.

Remarks

The EKFFT implementation was shown to require no knowledge about the excitation waveform and, if one can use solely the amplitude of the plasma response to extract the required information, it is a powerful method providing in parallel a frequency estimation of the magnetic perturbation. In particular, it is shown that not incorporating any knowledge about the frequency evolution of the excitation waveform can be made to have little impact in the amplitude estimation when compared to the one provided by the KFHE.

4.4.2 A real-time algorithm for the harmonic estimation and frequency tracking of dominant signal components in fusion plasma magnetic diagnostics

Introduction

This section extends the preliminary investigations of section 4.3 and compares the performance of three non-linear KF implementations in the FT configuration when applied to experimental data from magnetic diagnostics at JET. These are the EKFFT, robust Extended Kalman Filter Frequency Tracker (rEKFFT) and Unscented Kalman Filter Frequency Tracker (UKFFT).

When compared to other methods previously used in tokamak physics research for a similar kind of offline analysis (time-frequency analysis) such as the STDFT [Fig 06], the

wavelet transform [Poli 08], the Wigner-Ville distribution [Biza 99] or the Choi-Williams distribution [Figu 04], the KF based methods presented herein stand out by possessing a set of characteristics that make them particularly suitable for real-time applications. Whereas the other methods act on blocks of data, KF methods act on streams of data on a *point-by-point* basis. This means that, for the latter, the temporal resolution is the sampling period of the raw signals themselves while typically, in the case of the former, the temporal resolution is often order(s) of magnitude smaller than that. The main reason for this is that, in the case of KF methods, frequency is treated as a continuous quantity whereas other methods use a frequency grid whose resolution typically depends on the number of points actually processed. This means that a finer time resolution implies reprocessing the vast majority of samples already processed previously which has also the side effect of adding to the computational burden. KF methods on the other hand, because they model a system's state transition under the hypothesis of it being a Markov process, do not require the history of past measurements. Also, somewhat related to this are the issues of latency (crucial in the case of real-time feedback control) and dispersion in time localization. Whereas in the case of KF methods latency is small also because the estimates are typically very localized in the vicinity of the latest measurement, for the other methods latency is typically half the number of samples used in the analysis (which can be large especially when aiming at high frequency resolutions) and time localization is dispersed over the time span of processed samples.

Magnetic sensor signals

Magnetic diagnostics are used extensively in present fusion devices. These include, for instance, pick-up coils, saddle coils, Rogowsky coils and flux loops. A vast set of information can be inferred from magnetic measurements including the plasma current and the plasma boundary, both crucial for the real-time control of the experiment. The common underlying physical principle of these sensors' measurements is Faraday's law of induction where an induced voltage proportional to their number of turns and to variations of the magnetic flux across their cross-section is observed. At JET, and in other tokamaks in general, magnetic perturbations produced by MHD activity are identifiable as non-stationary signal components in pick-up coil signals located at the periphery of the plasma in the interior of the vacuum vessel. For the particular case of JET, for instance, the pick-up coils, of cylindrical shape, have typically an effective area of 0.064 m^2 for 70 turns of titanium wire with a 0.5 mm diameter. They were designed to measure within the bandwidth of 0-500 kHz.

In this work, the MHD frequencies of interest lie in the range of $\approx 0-40 \text{ kHz}$. Since the pick-up coil signals used in this work are typically digitized and made available for offline analysis at rates of 1 or 2 MSample/s, a decimation is performed in order to mimic a real-time system with a raw data acquisition rate of 200 kSample/s, thus relaxing the computational burden whilst guaranteeing adequate bandwidth in the spectral range of interest.

KF based frequency tracking

As shown in [Figu 06], the accuracy of amplitude estimations obtained by the short-time discrete Fourier transform of non-stationary signals is strongly dependent on the signal itself and on the number of points used in the sliding window of the analysis. This is a consequence of the signal's instantaneous frequency seldom coinciding with the frequency

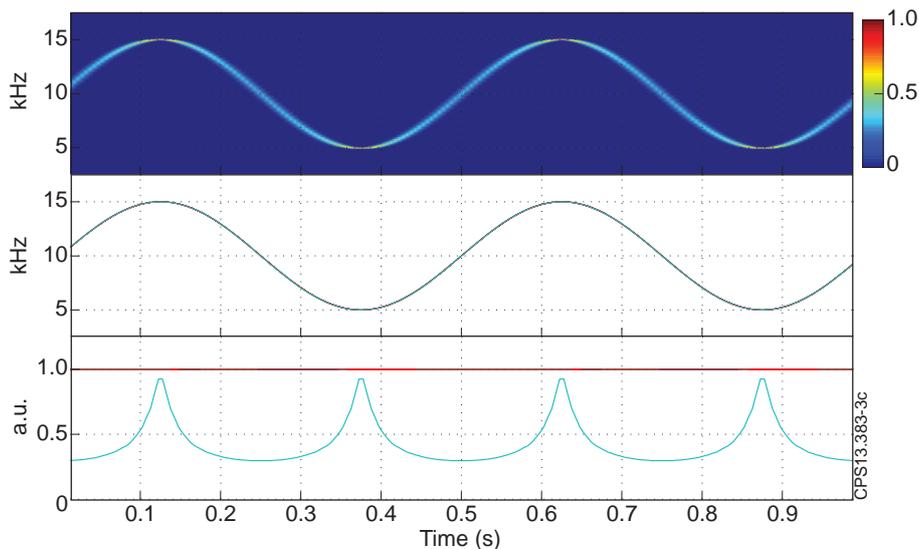


Figure 4.10: **Top plot:** raw signal spectrogram; **2nd plot:** estimated frequency; **Bottom plot:** estimated amplitude. **Blue:** EKFFT; **Green:** rEKFFT; **Red:** UKFFT; **Cyan:** DFT.

discretization bins, the spectral broadening effect and the trade-off of time and frequency resolutions combined with the dispersion in the time localization of the spectrogram. In principle, amplitude estimations obtained with non-linear KF-based methods in the FT configuration are immune to these effects. As an example, figure 4.10 shows a comparison of the frequency and amplitude estimations obtained with the EKFFT, rEKFFT and UKFFT methods and with the conventional spectrogram (STDFT). The original signal is an FM signal (carrier frequency oscillating between 5 and 15 kHz with a 2 Hz modulation) of unitary amplitude sampled at a rate of 200 kSample/s. The spectrogram is performed on sliding windows of 2000 samples (10 ms) with an overlap of 75%. It is clear from figure 4.10 that, although frequency estimations coincide (see middle plot), the spectrogram notoriously underestimates the true signal amplitude in regions of increased non-stationarity (bottom plot). On the contrary, all KF-based amplitude estimations overlap at the expected value of 1. Again, as shown in [Figur 06], better estimates can be achieved by using a smaller window at the cost of a poorer frequency resolution. Making the window too small introduces spurious oscillations thus degrading the quality of the estimates. Regarding the KF-based methods it is worth noting that, unless their bandwidth is intentionally trimmed by excessively privileging the theoretical model in detriment of the measurements, frequency and amplitude estimates will be accurate as shown in figure 4.10.

The following section will present the results obtained with the implementation of the FT configuration to the frequency tracking and harmonic estimation of MHD activity in JET plasma discharges.

Results

This section presents the results of applying some of the aforementioned non-linear KF based methods in the FT configuration to the simultaneous estimation of instantaneous frequency and amplitude of signal components in magnetic pick-up coils at JET. For this purpose, magnetic coil data for a set of pulses with identified MHD activity was used.

Results from both the iEKF and the soEKF are not presented herein. The former brings no added value with respect to the EKF in cases where the measurement equation is linear, which is the case of the FT configuration, see equation (4.3). The latter was found to be too sensitive to disturbances hence highly unstable.

Single-tone estimation and tracking

The first test was to estimate the instantaneous amplitude and frequency of the dominant component in a magnetic coil signal using the EKFFT, rEKFFT and UKFFT as described in the previous sections.

Both the EKFFT and rEKFFT have been configured with measurement variance, $R = 1$, and process covariance matrix of the form $Q = Q^{diag} = [q^a \ q^a \ q^f]$ with $q^a = 10^{-2}$ and $q^f = 10^{-4}$. Because in the case of the EKFFT (and rEKFFT) the Kalman gain depends solely on the ratio of the process and measurement covariances [Bitt 00], the measurement variance needs not be accurately determined and can be conveniently set to unity. The values of q^a and q^f , chosen from the range $10^{-8} - 10^2$ tested in order of magnitude steps, have been found to produce a stable algorithm and to provide a good compromise between noise attenuation and phase response. Furthermore, in the case of the rEKFFT, the *robustness* parameter was set to $\gamma = 500$. Smaller values of γ were found to, in some cases, trigger matrix conditioning problems affecting the inversions of equation (2.80) whereas larger values of γ tend to revert the rEKFFT to the EKFFT. The UKFFT was configured with $\alpha = 10^{-3}$, $\kappa = 0$ and $\beta = 2$ as suggested by Wan and Van Der Merwe [Wan 00] and with $R = 4 \times 10^{-3}$ and $Q = Q^{diag} = [q^a \ q^a \ q^f]$ with $q^a = 10^{-6}$ and $q^f = 10^{-7}$. The measurement variance, R , has been pre-calculated using experimental data and its typical value does not change from pulse to pulse. Once more, the values of q^a and q^f , chosen from the range $10^{-14} - 10^2$ tested in order of magnitude steps, have been found to produce a stable algorithm and to provide a good compromise between noise attenuation and phase response. Furthermore, all algorithms' initial state guess are based on the original signal's first available sample and a phase increment coherent with a 10 Hz frequency.

The main results are shown in figure 4.11 for JET pulse 68628. The top plot shows the spectrogram of the original signal using the STDFT. The two black lines delimit the area in the time-frequency plane containing the dominant component. This is used for the offline calculation of both the frequency and amplitude of the mode for comparison with the Real-Time (RT) processing algorithms. The second plot shows a comparison of the instantaneous frequency estimated by each algorithm. One can immediately see that, although the three RT methods provide accurate estimations most of the time, both the EKFFT and UKFFT lose track between 14.45 – 14.65 s and 13.95 – 14.42 s respectively. The third plot is the frequency estimation normalized error exhibiting the exact same behaviour of the previous plot and the fourth plot shows the instantaneous amplitude estimation. Here, again all RT methods agree amongst themselves and with the DFT most of the time except for the intervals where the EKFFT and UKFFT lose track. The fifth plot shows the amplitude estimation normalized error and the bottom plot shows the normalized residue r_k , given by equation (4.20) where z_k and $x_{1,k}$ are respectively the raw signal and the raw signal estimate at time index k . In the results presented herein, $N = 2000$ (equivalent to a 10 ms windowed calculation) was used. As expected, a noticeable increase of the normalized residue is observed in the regions where

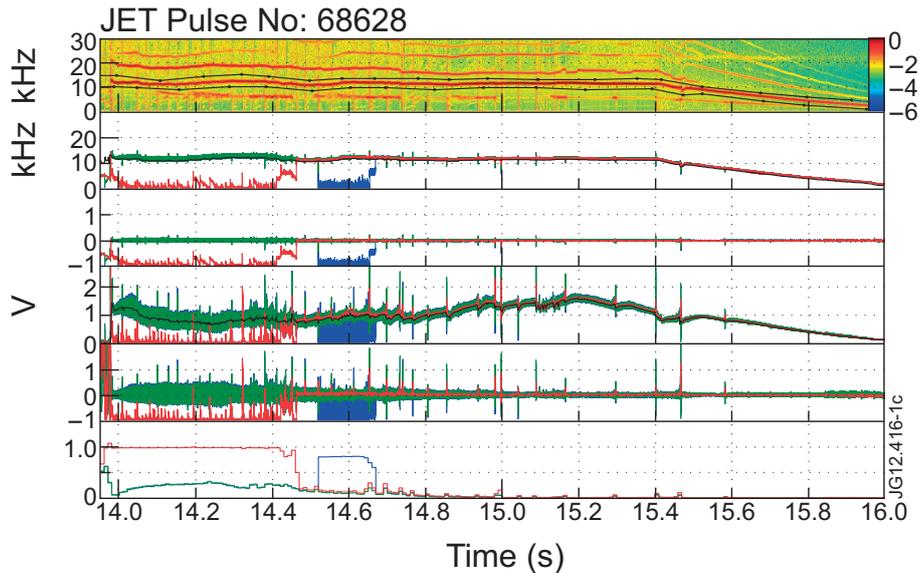


Figure 4.11: JET pulse 68628 - **Top plot**: raw signal spectrogram (\log_{10} scale); **2nd plot**: estimated frequency; **3rd plot**: normalized frequency estimation error; **4th plot**: estimated amplitude; **5th plot**: normalized amplitude estimation error; **Bottom plot**: 10 ms averaged residue. **Blue**: EKFFT; **Green**: rEKFFT; **Red**: UKFFT; **Black**: DFT.

the algorithms lose track.

$$r_k = \frac{\sum_{i=0}^{N-1} (z_{k-i} - x_{1,k-i})^2}{\sum_{i=0}^{N-1} z_{k-i}^2} \quad (4.20)$$

Three immediate conclusions can be drawn from this figure: 1- all RT methods are capable of providing accurate estimates; 2- the EKFFT and UKFFT are sometimes not robust enough and lose track of the frequency component; 3- the residue provides a measure of the estimation's accuracy.

Because of the spectral variety of typical magnetics signals in tokamak research, as exemplified in figure 4.11, it is natural that the signal component estimation provided by these RT algorithms is contaminated by other neighbouring and/or large components. In order to address this question, figure 4.12 shows the amplitude ratio in the time-frequency plane between the single-tone estimate, $x_{1,k}$, and the original signal, z_k . In order to obtain this figure, a STDFT with the exact same time-frequency resolution is applied to the original signal and the single-tone estimates obtained with each RT algorithm, thus making it possible to calculate all component amplitudes in the same time-frequency mesh. Apart from the already mentioned time regions where the EKFFT and UKFFT lose track it can immediately be seen that both the EKFFT and rEKFFT methods present similar levels of inter-frequency contamination whereas the UKFFT's performance is clearly superior. It can for example be seen that, for the strongest component with a frequency immediately above the dominant one and, at around 15s (~ 16 kHz) the UKFFT provides an attenuation of at least 60% while the others' never more than 30%. So at least a factor of two in that sense.

Recalling from (2.34) that the posterior state estimate contains a term where the

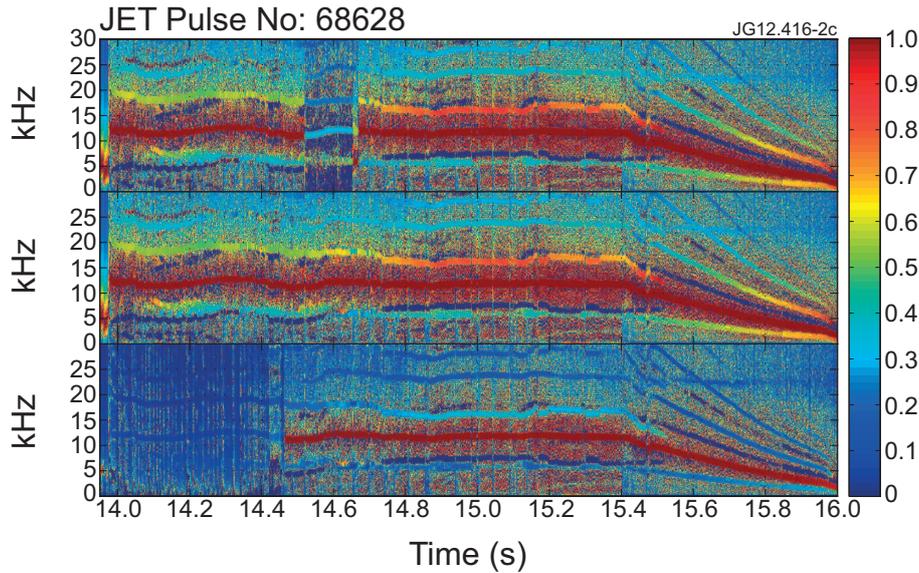


Figure 4.12: JET pulse 68628 - Normalized amplitude inter-frequency contamination of the estimated signal - **Top plot:** EKFFT; **Middle plot:** rEKFFT; **Bottom plot:** UKFFT.

KF gain is multiplied by the innovation ($\mathbf{z}_k - \mathbf{C}_k \hat{\mathbf{x}}_k$), it is clear that any signal component not being simultaneously estimated will be propagated into the final result.

So far it is clear that if the EKFFT or the UKFFT are to be used for this purpose, they will need to be more robust to unmodelled signal disturbances. Similarly, it is desirable to increase the contamination rejection performance of both the EKFFT and rEKFFT.

Adaptive single-tone estimation and tracking

As a result of the shortcomings of the RT methods identified previously, an adaptive scheme has been devised to increase both robustness and neighbouring frequencies' rejection. The idea is to instruct the filter to believe more in the model rather than in the measurement whenever the residue is low and the converse for larger residue values. Believing in the model means increasing the neighbouring frequency rejection and reducing the estimation noise at the expense of a lower resilience to unmodelled signal disturbances. As explained before, instructing the filter to trust the model more than the measurement is achieved with low Q/R values while the converse is achieved with high Q/R values.

The approach followed here has been to adapt the process covariance matrix according to the residue, r_k , by letting $q_k^f = 10^{-6+4r_k}$ and $q_k^a = 10^2 q_k^f$ in the cases of the EKFFT and rEKFFT. This is equivalent to adopting the configuration used in the non-adaptive case for $r_k = 0.5$ and increasing the rate of belief in the model for $0 < r_k < 0.5$ while doing the opposite for $0.5 < r_k < 1$. This can be verified by taking the limits $r_k \rightarrow 0$, in which case $q^f \rightarrow 10^{-6}$ and $q^a \rightarrow 10^{-4}$, and $r_k \rightarrow 1$, in which case $q^f \rightarrow 10^{-2}$ and $q^a \rightarrow 1$, and comparing it with the non-adaptive settings ($q^f = 10^{-4}$ and $q^a = 10^{-2}$).

As for the UKFFT, a parameter scan in order of magnitude steps had to be done as a straightforward approach equivalent to the one used for the EKFFT and rEKFFT resulted in poor and often unstable results. In this case $q_k^f = 10^{-8+6r_k}$ and $q_k^a = 10^3 q_k^f$ have been found to produce a stable algorithm and to provide a good compromise between noise attenuation and phase response over the range of possible residue values. The configuration parameters used in these adaptive relations have been demonstrated to be

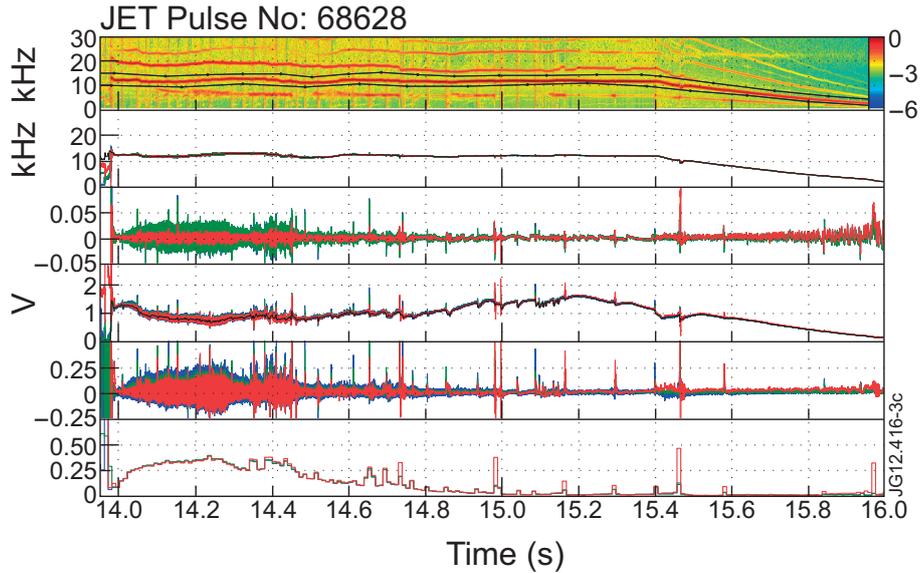


Figure 4.13: JET pulse 68628 - Adaptive estimation - **Top plot**: raw signal spectrogram (\log_{10} scale); **2nd plot**: estimated frequency; **3rd plot**: normalized frequency estimation error; **4th plot**: estimated amplitude; **5th plot**: normalized amplitude estimation error; **Bottom plot**: 10 ms averaged residue. **Blue**: EKFFT; **Green**: rEKFFT; **Red**: UKFFT; **Black**: DFT.

robust enough to be used in a wide variety of JET pulses where a clearly dominant signal component (at least a factor of 2 larger in amplitude than the next strongest one, to be on the safe side) is present.

The benefits of this adaptive scheme are visible in Figs. figure 4.13 and figure 4.14. Not only the EKFFT and UKFFT are now capable of tracking the signal tone at all times, also the neighbouring component rejection performance is now comparable amongst RT algorithms. Still, the frequency and amplitude estimation errors are smaller for the UKFFT when compared with the other methods. It can also be seen that the estimation performance is worst roughly between 14 and 14.5s. Even in these conditions worst case frequency and amplitude estimates are achieved with errors below 5% and 25%, respectively. After that, because the ratio between the dominant (the tone being estimated) and the next strongest tone increases, frequency and amplitude estimates exhibit errors below 2% and 10%, respectively.

Another example is shown in Figs. figure 4.15 and figure 4.16 this time for JET pulse 67694. It is clear that the proximity of intermittent lower frequency activity (known as sawtooth activity) between 14.5 and 16s deteriorates the estimates. Subsequently, as the distance between the dominant mode and this lower frequency activity increases, the estimates improve and it can be seen that the frequency estimation error is below 2% and the amplitude estimation error is below 20%.

Another example is shown in Figs. figure 4.17 and figure 4.18 for JET pulse 79743. Frequency estimation error is of the order or below 1% and, for the amplitude estimation error, this value is roughly 20%. The effect of sawtooth activity in the 20-23s range is evident in the residue although never causing the filters to stop tracking the dominant component. Also, figure 4.18 shows that the contamination rejection performance of all algorithms is comparable.

This section has demonstrated robust and reliable tracking along with frequency and

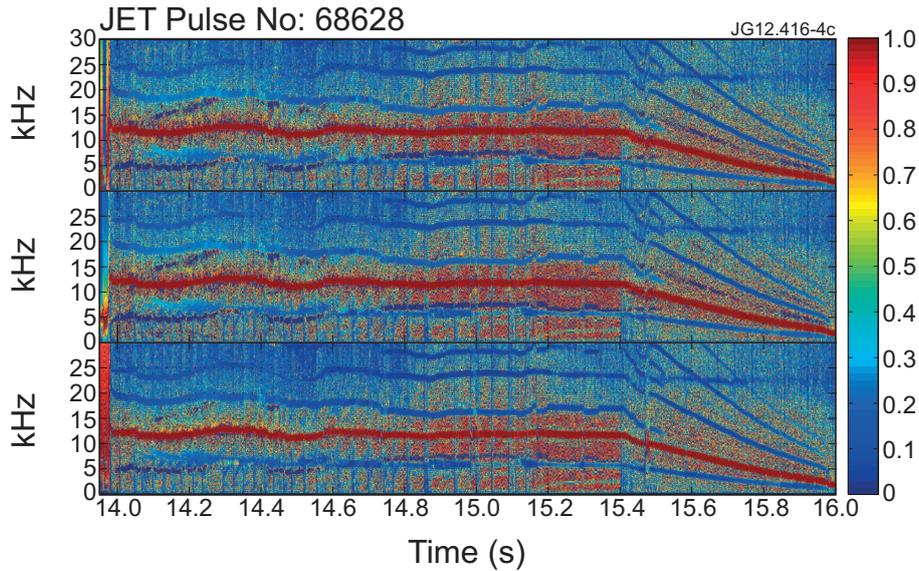


Figure 4.14: JET pulse 68628 - Adaptive estimation - Normalized amplitude inter-frequency contamination of the estimated signal - **Top plot:** EKFFT; **Middle plot:** rEKFFT; **Bottom plot:** UKFFT.

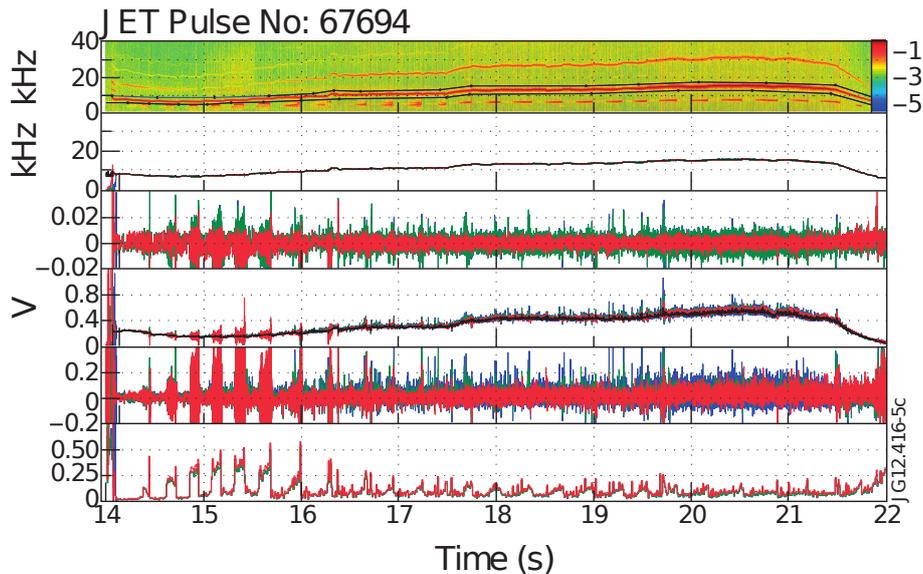


Figure 4.15: JET pulse 67694 - Adaptive estimation - **Top plot:** raw signal spectrogram (\log_{10} scale); **2nd plot:** estimated frequency; **3rd plot:** normalized frequency estimation error; **4th plot:** estimated amplitude; **5th plot:** normalized amplitude estimation error; **Bottom plot:** 10 ms averaged residue. **Blue:** EKFFT; **Green:** rEKFFT; **Red:** UKFFT; **Black:** DFT.

amplitude estimations of dominant MHD activity in JET magnetic signals. The normalized residue has been used to adjust the process covariance by balancing the noise rejection and inter component contamination with the disturbance resilience accordingly. Also, it has been shown that the residue can be used simultaneously as a measure of the quality of the estimations thus providing a critical real-time validation check if some kind of feedback control action is to be devised. Furthermore, one will note the fairly regular transients observed especially in the spectrograms and amplitude estimates of pulse 68628.

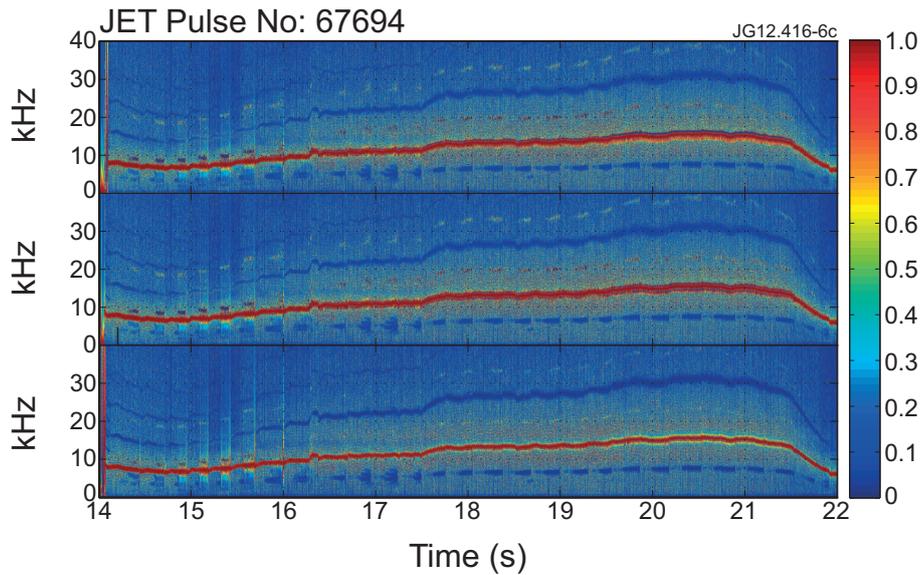


Figure 4.16: JET pulse 67694 - Adaptive estimation - Normalized amplitude inter-frequency contamination of the estimated signal - **Top plot:** EKFFT; **Middle plot:** rEKFFT; **Bottom plot:** UKFFT.

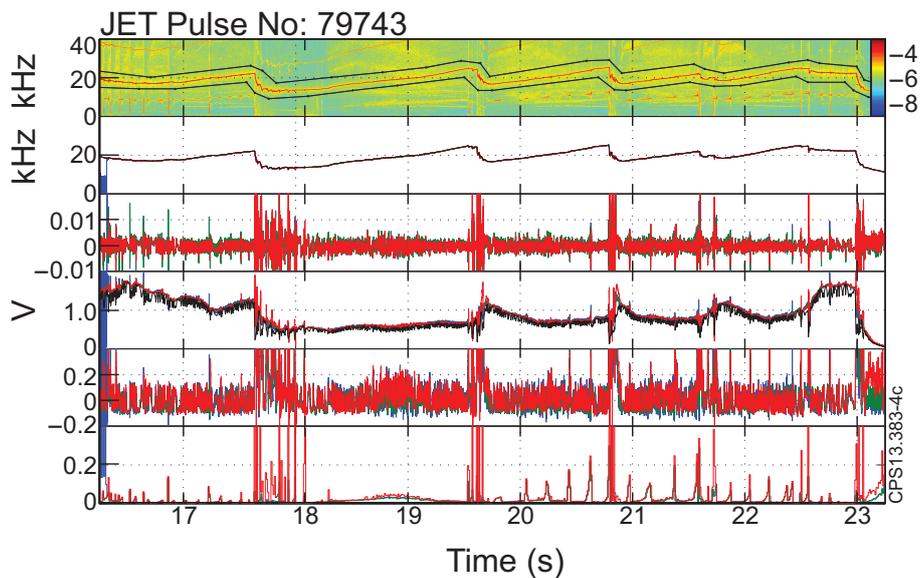


Figure 4.17: JET pulse 79743 - Adaptive estimation - **Top plot:** raw signal spectrogram (\log_{10} scale); **2nd plot:** estimated frequency; **3rd plot:** normalized frequency estimation error; **4th plot:** estimated amplitude; **5th plot:** normalized amplitude estimation error; **Bottom plot:** 10 ms averaged residue. **Blue:** EKFFT; **Green:** rEKFFT; **Red:** UKFFT; **Black:** DFT.

These are originated by events known in tokamak research as ELM [Kami 07]. Although these pulses exhibit moderately large and regular unmodelled ELM disturbances, all filter algorithms remain on track of the dominant component thus demonstrating to be effectively resilient to these disturbances.

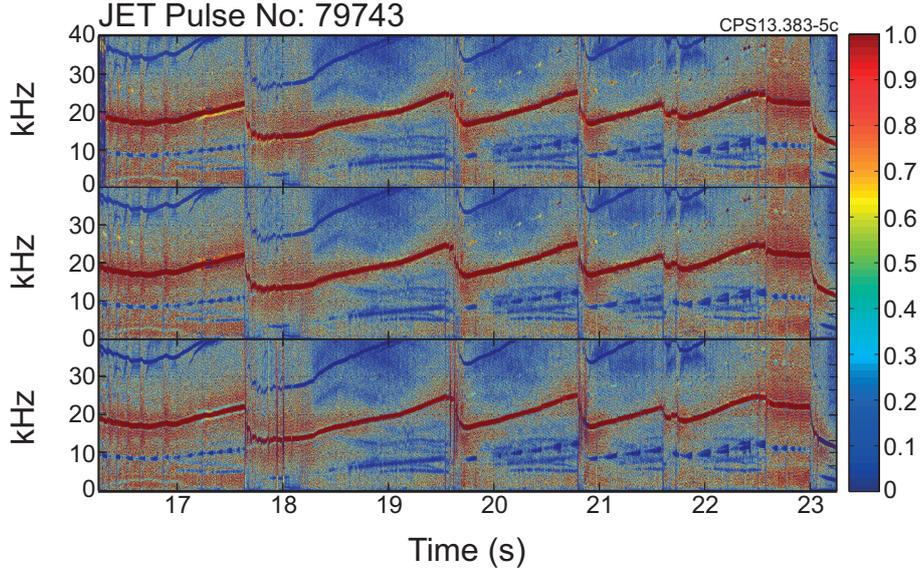


Figure 4.18: JET pulse 79743 - Adaptive estimation - Normalized amplitude inter-frequency contamination of the estimated signal - **Top plot:** EKFFFT; **Middle plot:** rEKFFFT; **Bottom plot:** UKFFT.

Adaptive multi-tone estimation and tracking

As previously mentioned, magnetic signals of tokamak plasma discharges can contain rich spectral sets and are therefore not restricted to a single frequency component. Examples of such spectral diversity are evident in all spectrograms presented herein. Whereas the previous section demonstrated robust tracking of the dominant frequency component in a magnetic signal, this section illustrates results obtained with a multi-tone tracking mechanism based on the same fundamental principles and algorithms. This is basically achieved by extending the frequency tracking configuration presented in section 4.4.2 to accommodate simultaneous multi-tone estimation. For practical and symmetry reasons which will be explained later on, this work will address simultaneous three tone estimation. Expressing the model equation for the frequency tracker in the usual state space notation yields $\hat{x}_k = M_k x_{k-1}$ where the block diagonal process matrix, M_k , is now given by equation (4.21). The $M_{j,k}$ matrices, equation (4.22), are immediately recognized as the fundamental blocks responsible for the estimation of each independent signal component and the measurement relation is now given by equation (4.23). The process model is essentially a *three-independent-tone* harmonic estimator and frequency tracker and the measurement equation reflects the assumption that the raw signal is described as the sum of these three dominant components plus normally distributed noise.

$$M_k = \begin{bmatrix} M_{1,k} & & \\ & M_{2,k} & \\ & & M_{3,k} \end{bmatrix} \quad (4.21)$$

$$M_{j,k} = \begin{bmatrix} \cos(x_{3j,k-1}) & -\sin(x_{3j,k-1}) & 0 \\ \sin(x_{3j,k-1}) & \cos(x_{3j,k-1}) & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad j = 1, 2, 3 \quad (4.22)$$

$$z_k = [1 \ 0 \ 0 \ 1 \ 0 \ 0 \ 1 \ 0 \ 0] [\hat{x}_{1,k} \ \hat{x}_{2,k} \ \hat{x}_{3,k} \ \hat{x}_{4,k} \ \hat{x}_{5,k} \ \hat{x}_{6,k} \ \hat{x}_{7,k} \ \hat{x}_{8,k} \ \hat{x}_{9,k}]^T \quad (4.23)$$

The adaptive mechanism used to adjust the process covariance matrix is the same as the one used in the previous section, both the expressions for q_k^f and q_k^a have been preserved, only now $Q_k = Q_k^{diag} = [q_k^a \ q_k^a \ q_k^f \ q_k^a \ q_k^a \ q_k^f \ q_k^a \ q_k^a \ q_k^f]$.

With the inclusion of extra degrees of freedom in the process model there is an increased likelihood that the filter will converge to a local minimum instead of the desired solution. One of the main problems is that, for instance in the presence of a strictly dominant frequency component, and for symmetry reasons, all independent estimates will (despite being fed with different initial guesses) eventually converge to the strongest tone. Such a situation would not only make the three independent tone estimate redundant, but would also complicate its interpretation as each estimate would account for an equal share of the dominant signal component. Therefore, unlike the single-tone estimation, three constraints need to be imposed in order to ensure consistent operation of the filters in these circumstances. The first constraint is to enforce a minimum inter-frequency separation amongst estimated components. The second constraint is to enforce a predefined estimate (associated with either $M_{1,k}$, $M_{2,k}$ or $M_{3,k}$) to track the strongest signal component and the third constraint is to enforce each of the remaining two non-dominant estimates to lie respectively (frequency-wise) above and below the largest component. The first constraint will stop the independent estimates from converging to the same tone while the second and third constraints define a simple priority rule when imposing the inter-frequency separation. Indeed, a more generic priority rule based on the sorting of each tone's instantaneous amplitude estimation in descending order of magnitude could be devised so that both the second and third constraints would no longer be required. However, let one consider for example the situation where the estimated component with the smallest amplitude lies, frequency-wise, between the other two estimates. In that case, if the frequencies of the two largest tones approach, the estimate of the weakest component might be left without a frequency range to *exist* thus forcing a decision to be made on whether it should be moved to the highest or the lowest frequency range of all estimates. Dealing with these sort of situations not only greatly adds up to the complexity of the algorithm but is also irrelevant for the proof of principle aimed herein.

Therefore, whenever the frequency separation between the dominant and a non-dominant estimate is found to be below the predefined minimum value, the posterior estimate is altered by enforcing the phase increment of the non-dominant tone to be in a valid region. The convention used herein is to let $M_{2,k}$ be associated with the dominant signal component and allow $M_{1,k}$ and $M_{3,k}$, respectively, to lie below and above the dominant tone's estimated frequency.

Results obtained from applying this *three-independent-tone* simultaneous estimation mechanism to JET pulse 68628 are shown in Figs. (figure 4.19), (figure 4.20) and (figure 4.21) respectively for the EKFFT, rEKFFT and UKFFT algorithms. Minimum inter-frequency separation is 5 kHz and estimates < 5 kHz for the dominant tone are not allowed. The top plot in these figures is a spectrogram of the original magnetic signal in the time-frequency region of interest. The second plot shows the frequency estimations obtained for each of the three components. The third, fourth and fifth plots show, respectively, the amplitude estimations of signal components a), b) and c) and the sixth plot shows the individual residues, calculated using (4.24), for each independent component

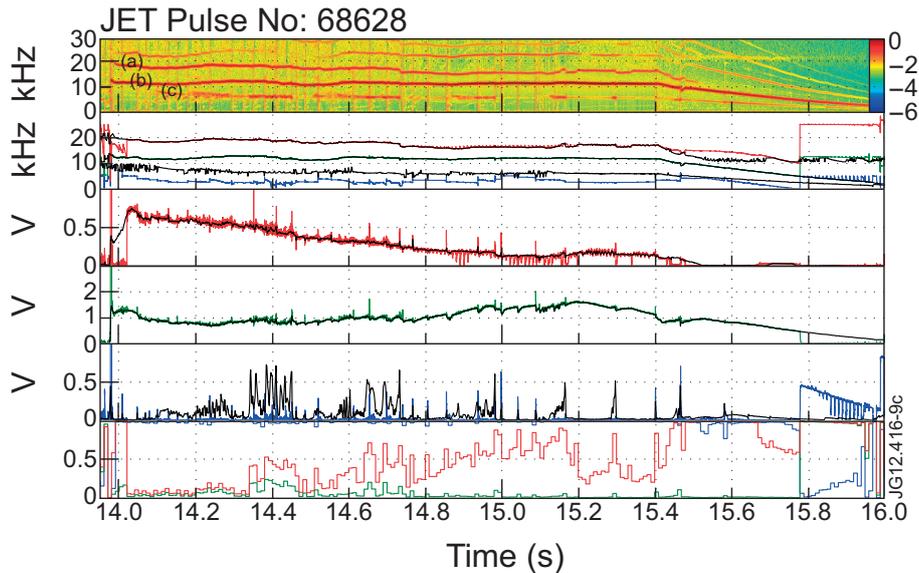


Figure 4.19: JET pulse 68628 - Three tone adaptive EKFFT - **Top plot**: raw signal spectrogram (\log_{10} scale); **2nd plot**: estimated frequencies; **3rd plot**: amplitude of tone a); **4th plot**: amplitude of tone b); **5th plot**: amplitude of tone c); **Bottom plot**: residues. **Red**: tone a) estimate; **Green**: tone b) estimate; **Blue**: tone c) estimate; **Black**: DFT.

estimate, j , of a total of $M = 3$ estimates.

$$r_{j,k} = \frac{\sum_{i=0}^{N-1} \left(z_{k-i} - \sum_{l=1}^M x_{3(l-1)+1,k-i} \right)^2}{\sum_{i=0}^{N-1} \left(z_{k-i} - \sum_{l=1, l \neq j}^M x_{3(l-1)+1,k-i} \right)^2}, \quad j = 1, 2, \dots, M \quad (4.24)$$

The first remark is that the estimation of the dominant signal component, tone b), is unaffected when compared to the single-tone adaptive estimation presented in section 4.4.2. This is clearly a crucial requirement given the importance of the dominant tone. The second remark is that the estimation of the non-dominant signal components is not as good as the dominant one. This is a consequence of their lower amplitude and therefore lower signal-to-noise ratio and increased vulnerability to inter-frequency contamination. For instance, in the case of tone a), whereas the EKFFT starts tracking it after a short transient, the rEKFFT starts by tracking the signal component immediately above it and when this one's strength starts to dim, moves on to track tone a). The UKFFT starts tracking tone a) immediately. As for tone c), the EKFFT was unable to track it whilst both the rEKFFT and the UKFFT present similar tracking capabilities frequency and amplitude-wise. The third remark, as crucial as the first one, is that the quality of the individual estimates is evident and fully correlated with the corresponding residue.

The same type of results are shown in Figs. (figure 4.22), (figure 4.23) and (figure 4.24) for JET pulse 68558. In this case, all filter algorithms track the dominant component, b), seamlessly. The EKFFT was unable to track tone a) and seldom able to track tone c). On the other hand both the rEKFFT and UKFFT exhibit similar tracking capabilities regarding the non-dominant signal components.

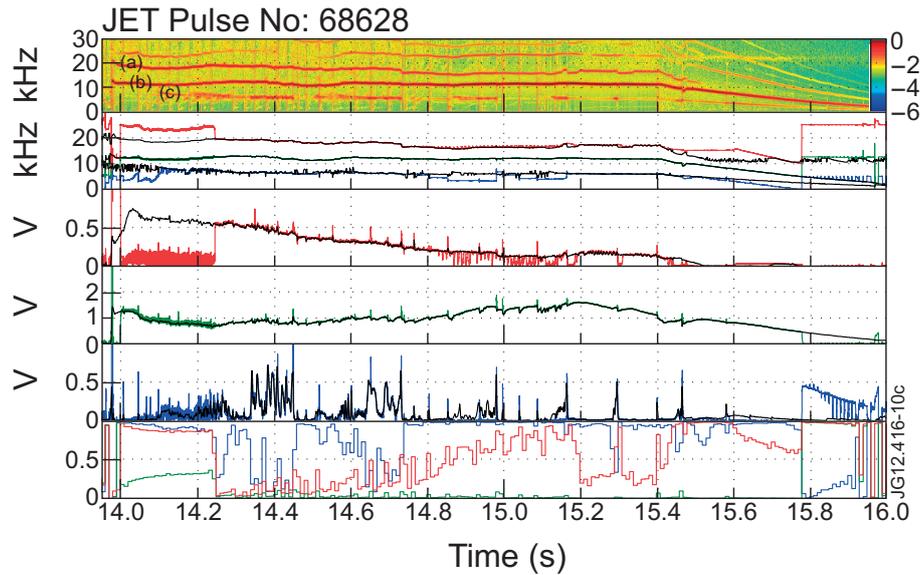


Figure 4.20: JET pulse 68628 - Three tone adaptive rEKFFT - **Top plot**: raw signal spectrogram (\log_{10} scale); **2nd plot**: estimated frequencies; **3rd plot**: amplitude of tone a); **4th plot**: amplitude of tone b); **5th plot**: amplitude of tone c); **Bottom plot**: residues. **Red**: tone a) estimate; **Green**: tone b) estimate; **Blue**: tone c) estimate; **Black**: DFT.

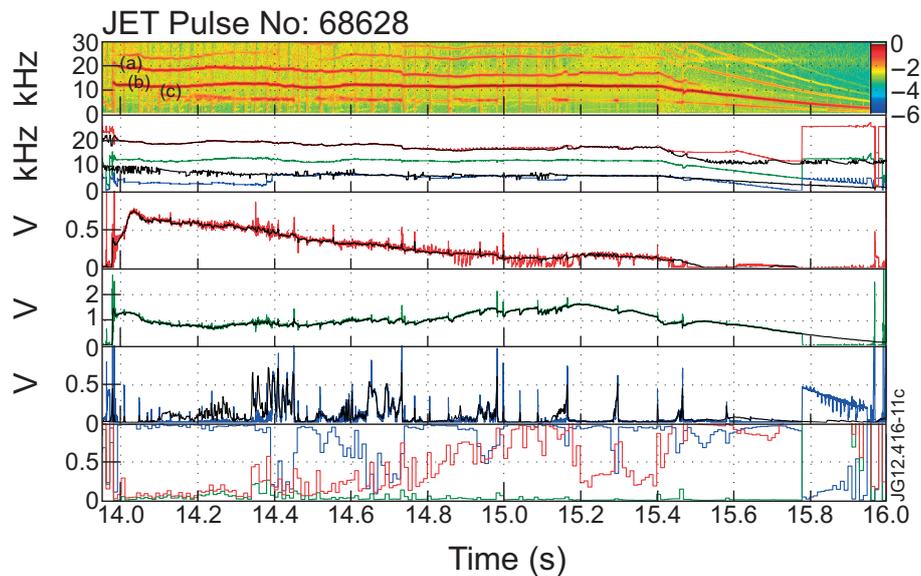


Figure 4.21: JET pulse 68628 - Three tone adaptive UKFFT - **Top plot**: raw signal spectrogram (\log_{10} scale); **2nd plot**: estimated frequencies; **3rd plot**: amplitude of tone a); **4th plot**: amplitude of tone b); **5th plot**: amplitude of tone c); **Bottom plot**: residues. **Red**: tone a) estimate; **Green**: tone b) estimate; **Blue**: tone c) estimate; **Black**: DFT.

This section demonstrates that simultaneous adaptive multi-component estimation is possible and does not seem to affect the dominant tone estimation. Although reliable estimates for the non-dominant components cannot be guaranteed at all times, the individual residue is always able to provide a measure of whether a particular independent tone estimation is actually tracking a meaningful signal component or not, again, crucial if part of a real-time feedback mechanism. Unlike the single-tone estimation case though,

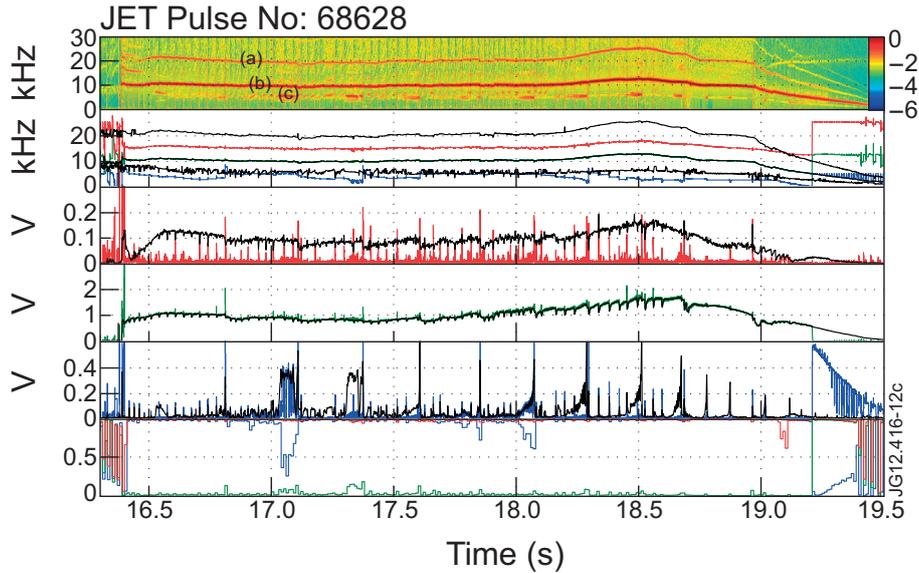


Figure 4.22: JET pulse 68558 - Three tone adaptive EKFFT - **Top plot**: raw signal spectrogram (\log_{10} scale); **2nd plot**: estimated frequencies; **3rd plot**: amplitude of tone a); **4th plot**: amplitude of tone b); **5th plot**: amplitude of tone c); **Bottom plot**: residues. **Red**: tone a) estimate; **Green**: tone b) estimate; **Blue**: tone c) estimate; **Black**: DFT.

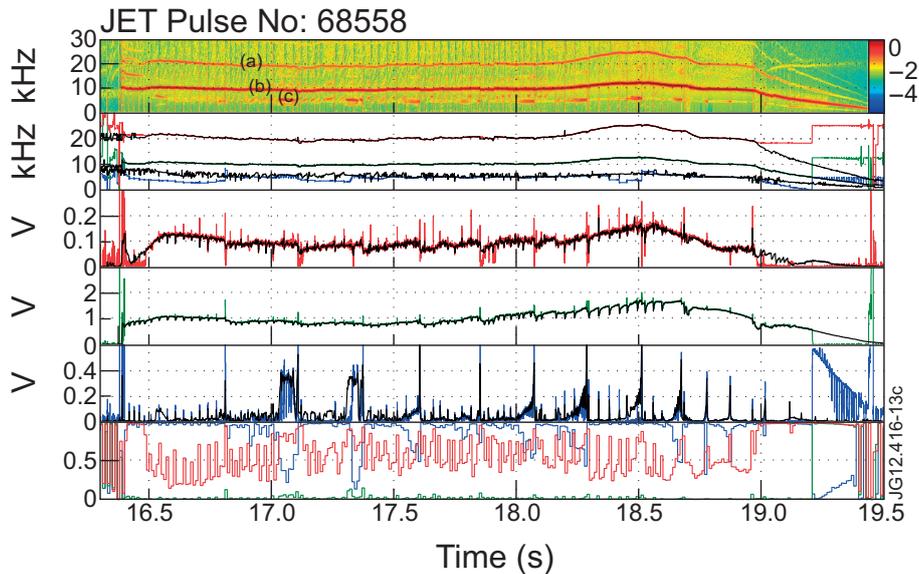


Figure 4.23: JET pulse 68558 - Three tone adaptive rEKFFT - **Top plot**: raw signal spectrogram (\log_{10} scale); **2nd plot**: estimated frequencies; **3rd plot**: amplitude of tone a); **4th plot**: amplitude of tone b); **5th plot**: amplitude of tone c); **Bottom plot**: residues. **Red**: tone a) estimate; **Green**: tone b) estimate; **Blue**: tone c) estimate; **Black**: DFT.

some constraints need to be established to ensure consistent operation of these types of filter. Furthermore, simultaneous multi-tone estimation introduces additional complexity and extra demand in terms of processing power when compared to the adaptive single-component estimation algorithm presented earlier. As an example, in the case of the EKFFT, the core filter algorithm requires 274 floating point operations (including two trigonometric function evaluations) in single-tone estimation whereas for the three-tone

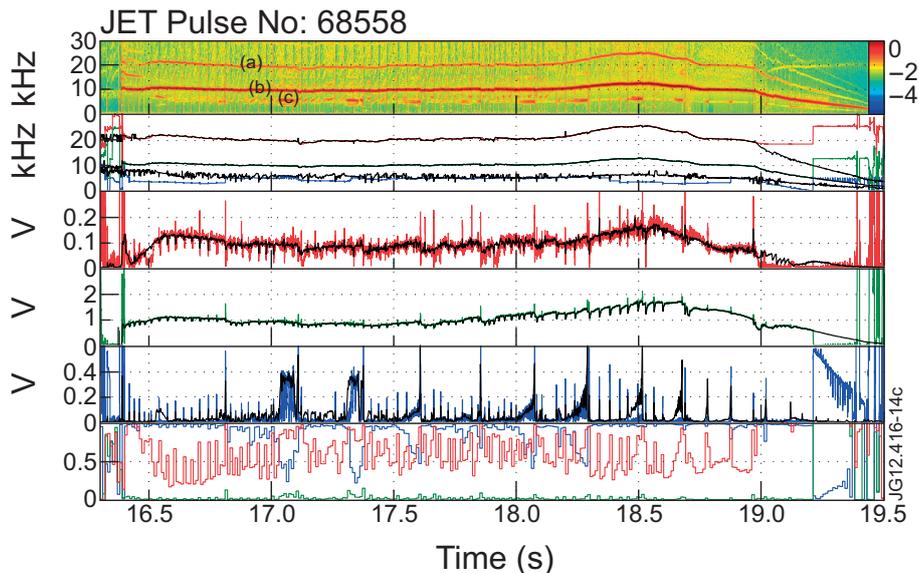


Figure 4.24: JET pulse 68558 - Three tone adaptive UKFFT - **Top plot**: raw signal spectrogram (\log_{10} scale); **2nd plot**: estimated frequencies; **3rd plot**: amplitude of tone a); **4th plot**: amplitude of tone b); **5th plot**: amplitude of tone c); **Bottom plot**: residues. **Red**: tone a) estimate; **Green**: tone b) estimate; **Blue**: tone c) estimate; **Black**: DFT.

estimation it requires 6272 floating point operations (including six trigonometric function evaluations) per cycle. This translates into a factor > 20 in terms of processing power requirements.

The next section shows how a single-tone non-linear adaptive KF FT can be incorporated in a general scheme for the real-time frequency and mode number analysis of MHD activity at JET.

4.4.3 Kalman filter methods for real-time frequency and mode number estimation of MHD activity in tokamak plasmas

Introduction

The successful operation of a magnetically confined fusion plasma power plant faces numerous challenges on various topical areas such as material science (e.g neutron resistant materials), breeding blanket technology, high performance plasma scenarios, active control of the plasma profiles/stability and confinement degradation/disruption mitigation. The latter three areas are closely interlinked given that it is extremely rare that a high performance plasma regime is not hampered by the onset and growth of detrimental plasma instabilities that, under certain cases, may lead to disruptions. Unless some active control of particular plasma instabilities is set, the figures of merit of the plasma performance, e.g. stored energy, plasma beta (ratio of plasma pressure to magnetic pressure), pressure gradients are inevitably limited and the tolerable heat loads to the plasma facing components of the device are eventually overcome. Active control mechanisms for the sawtooth instability [Saut 02, Chap 12], Neoclassical Tearing Modes (NTMs) [H Zo 07, Haye 06], resistive wall modes [In 06b, Liu 00, Baru 12] and ELM [Evan 08, Lian 07] are therefore of pivotal importance. Such control systems will ultimately be complemented with other systems targeting the plasma performance to address for instance the current density

control and the pressure profile control in the presence of internal transport barriers. Additionally, the development of successful feedback stabilization schemes for plasma instabilities that increase the core plasma radial transport such as the sawtooth and for potentially disruptive instabilities such as (2,1) NTMs should be considered in terms of integrated control implementations since, as in other cases, the dynamics of these two instabilities is strongly coupled. In fact, the internal reconnection instability localized near the rational magnetic surface $q = 1$, associated with the sawtooth oscillations [Hast 97], is a potential drive, for sufficiently long sawtooth periods [Butt 03, Butt 04], for the generation of *seed islands* at other rational $q = m/n$ (m, n integers) magnetic surfaces in the plasma. For sufficiently large seed islands, the local radial plasma transport is increased and the pressure profile flattened. This in turn drives the NTM unstable since, as is well known, the NTM instability is driven by the local reduction of the bootstrap current (proportional to the pressure gradient) [Hegn 98]. Feedback stabilization schemes for macroscopic plasma instabilities are inherently more robust and efficient when a more complete characterization of the instability is provided. This entails providing information on the temporal evolution of the mode, e.g. propagating frequency in the lab frame and growth rate but also on the spatial properties of the mode, e.g. topology of the associated magnetic structure (kink structure or island structure with *O-points* and *X-points*) and, for internal modes, the rational flux surface $q(\psi_{pol}) = m/n$ where the mode is resonant. Here, ψ_{pol} stands for the poloidal magnetic flux (acting as a *radial coordinate* in the magnetic flux coordinate system) and m and n stand for the poloidal and toroidal mode numbers respectively. For the particular case of NTMs, for instance, it is essential to know the radial location of the mode when developing feedback stabilization schemes based on Electron Cyclotron Resonance Heating (ECRH) and Electron Cyclotron Current Drive (ECCD) [H Zo 07, Haye 06, Saut 04] so that the heating and current drive are precisely sourced within the island region. More accurately, in order to maximize the efficiency of the stabilization scheme, the EC deposition width should be smaller than the island width and the EC power modulated synchronously to the island's propagating phase in order to drive current primarily at the island's *O-point*. It is therefore reasonable to assume that, in this context, real-time tracking methods that extract the frequency, growth rate, mode numbers and phase of plasma instabilities are a valuable tool that assists the development and operation of plasma control schemes. In this work, a real-time method for the identification of the mode numbers of plasma perturbations as well as their amplitude and phase is derived, based on KF techniques. In this particular field of research, KF based methods have already been successfully used in the context of feedback stabilization of MHD activity, such as resistive wall modes, in both the HBT-EP [Hans 09] and DIII-D [In 06c] tokamaks. In the former it is used for estimating the in-phase and quadrature components of the poloidal field due to a rotating and growing $n=1$ mode and in the latter it is used mainly for discriminating ELMs from true resistive wall modes thus assisting in the mitigation of undesirable and inopportune control actions. In both applications, the KF has demonstrated its superior performance in terms of noise rejection while providing low latency estimates, typically with a throughput rate equal to the sampling frequency of the measurements.

The remainder of this section is organized as follows: first an introduction provides the fundamentals of mode number analysis and exposes the basic underlying principles of popular methods used for such purpose, then the proposed method is described and subsequently projected onto JET's typical plasma conditions and sensors' capabilities, availabilities and spatial distribution establishing a suitable combination of sensors by

pondering processing demands and estimation performance. Finally results are presented by mimicking the proposed method's performance under real-time conditions using historical sensor data from JET pulses.

Mode number analysis: a summary of conventional methods

When characterizing plasma instabilities, one easily distinguishes between temporal and spatial characteristics of the associated plasma perturbations in, e.g., magnetic field, plasma density, temperature and velocity. The former addresses the propagating frequency of the perturbations and their growth rate while the latter addresses the radial profile of the plasma perturbations and their toroidal and poloidal structure. Conventionally, a generic plasma perturbation is identified by a functional $f(t, r, \theta, \phi)$ where r is a flux surface labelling coordinate (a *radial coordinate* in the torus) and θ and ϕ are, respectively, poloidal and toroidal angles. An expansion in Fourier series for the (θ, ϕ) dependency yields (4.25) where $d\varphi/dt = \omega(t)$ represents the instantaneous mode frequency and m and n are, respectively, the poloidal and toroidal mode numbers.

$$f(t, r, \theta, \phi) = \text{Re} \left[\sum_n \sum_m f_{(m,n)}(t, r) e^{j(\varphi(t) + m\theta - n\phi)} \right] \quad (4.25)$$

Certain types of instabilities exhibiting a clear poloidal and toroidal symmetry, namely kink and tearing instabilities, are resonant at magnetic surfaces where the safety factor $q(\psi_{pol}) = d\psi_{tor}/d\psi_{pol}$ is rational (ψ_{tor} and ψ_{pol} are, respectively, the toroidal and poloidal flux) and as such the experimental identification of the poloidal (m) and toroidal (n) periodicity of the perturbation provides the signature of the rational $q(\psi_{pol}) = m/n$ surface where the mode is located. An illustration of a toroidal equilibrium evidencing a $q = 3/2$ magnetic surface (inner meshed surface), the topology of the helical magnetic field line at that surface and an array of three diagnostic sensors lying on an outer mesh (vis behind the first wall) is shown in figure 4.25.

The expansion of equation (4.25) is useful to decompose plasma perturbations into eigenmodes with eigenvalues given by the triplet (ω, m, n) . While the decomposition in the toroidal eigenvalues is easily justified when considering the equilibrium toroidally axisymmetric, the decomposition in poloidal eigenvalues is not as straightforward owing to the coupling between different poloidal harmonics. This coupling stems from the toroidicity and shaping of the plasma equilibrium and couples neighbouring poloidal harmonics already during the linear regime (exponential growth) of the perturbations. Ultimately, for each toroidal mode number n there will be a spectra of poloidal mode numbers, possibly resonant in some rational surfaces, with competing drive/damping mechanism contributing to the overall growth rate of the mode. For large aspect ratio (R_0/a) plasma equilibria, with R_0 and a the major and minor radius of the equilibrium, the coupling between poloidal harmonics is strongly reduced and the stability of (m, n) modes can be treated independently.

The fundamental relation for mode number analysis is given by (4.26) for the case of a single non-degenerate mode observed at angular frequency ω in the laboratory frame and with m and n , respectively, the poloidal and toroidal periodicities. It basically establishes the relation between the phase difference $\Delta\Psi_{ij}$ observed in the signal of two sensors (i and j), at angular frequency ω , separated $\Delta\theta_{ij}$ in the poloidal direction and $\Delta\phi_{ij}$

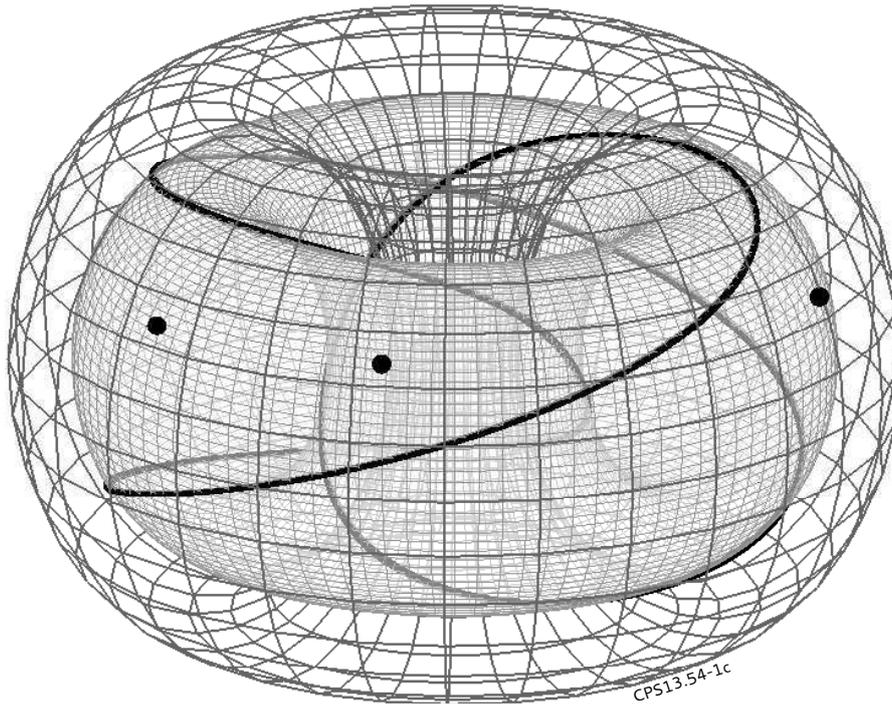


Figure 4.25: Example of a $q = m/n = 3/2$ flux surface and field line

in the toroidal direction.

$$\Delta\Psi_{ij}(\omega) = m(\omega)\Delta\theta_{ij} - n(\omega)\Delta\phi_{ij} \quad (4.26)$$

In this work the experimental sensors used are located at the same poloidal angle, deliberately privileging the determination of the toroidal mode number of plasma fluctuations. In practical terms, such a setup cancels out the term in (4.26) proportional to the poloidal angle variation between sensors. Equivalently, if for instance all sensors are located within the same poloidal cross-section, the term in (4.26) proportional to the toroidal angle separation vanishes. The choice of such a *toroidal sensor setup* arises from the previously mentioned fact that, while the determination of the toroidal mode number of plasma fluctuations in conventional tokamak geometries is a well-posed problem (excluding the non-linear evolution, a fluctuation can be characterized by a single toroidal mode number), the same does not hold for the determination of the poloidal mode number.

Three fundamental types of methods are commonly used for mode number analysis: Fourier, Singular Value Decomposition (SVD) and periodogram based methods. Within these types numerous variants, including hybrid ones, coexist. Methods can be characterized/benchmarked using a number of factors such as but not restricted to: a) whether they work in the time or frequency domain; b) whether they provide mode number, frequency, amplitude and/or phase information; c) time and frequency resolutions; d) dispersion in time localisation; e) capability of resolving mode degeneracy; f) whether sensor calibration is required; g) whether uniform sampling is required in time and/or space; h) resilience to disturbances. In this section, a basic introduction to the fundamentals of each method is given and, without loss of generality, the sensors probing the fluctuations are assumed to be magnetic flux pick-up coils for which measurements are acquired synchronously. The acquired signal is proportional to the time variation of the flux of the total magnetic field

(equilibrium plus fluctuations) through the coil's cross-section.

Spatial Discrete Fourier Transform Arguably, one of the simplest methods that can be automated to calculate the spatial periodicity of a magnetic perturbation using pick-up coil signals is by means of a Spatial Discrete Fourier Transform (SDFT). Assuming a set of N coils, the idea is merely to take a snapshot of each coil measurement $s(t, \phi)$, at time t and toroidal position ϕ , and evaluate its periodicity n in the toroidal direction using a discretised version of (4.27) where the integral is replaced by the summation over the discrete toroidal positions ϕ_i of the array of coils and $S(t, n)$ is the complex-valued Fourier component of n at time t .

$$S(t, n) = \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} s(t, \phi) e^{-jn\phi} d\phi, \quad i = 1, 2, \dots, N \quad (4.27)$$

The method's temporal resolution is the sensor signal's sampling period. It can provide the amplitude and phase of a pre-defined set of n values but it does not provide any information nor discrimination of frequency components. Furthermore, usage of the popular FFT [Cool 65] algorithm to efficiently compute the DFT requires spatial uniform sampling which is seldom available in experimental devices due to mechanical and technological restrictions. In addition, the accuracy of the analysis is very much dependent on the total number of sensors available.

Coherence spectrum The coherence spectrum [Rabi 75], sometimes referred to as cross-power spectral density is another Fourier based method commonly used for the mode number calculation of MagnetoHydroDynamic (MHD) activity in magnetically confined fusion plasmas [Ferr 92]. Considering two pick-up coil signals $x(t)$ and $y(t)$ with complex-valued Fourier transforms, $X(\omega)$ and $Y(\omega)$, respectively, evaluated at angular frequency ω . Then, if Y^* denotes the complex conjugate of Y , X^* the complex conjugate of X and $\langle \cdot \rangle$ represents the ensemble average [Gned 06], equation (4.28) gives the complex-valued cross power spectral density $\chi(\omega)$ of both signals. Since determining the ensemble average is not feasible when dealing with unique time series of data, and assuming that the signals are stationary in a certain interval and that the process is reasonably ergodic, the ensemble average can be replaced by a suitable time average.

$$\chi(\omega) = \frac{\langle X(\omega)Y^*(\omega) \rangle}{\langle X(\omega)X^*(\omega) \rangle^{\frac{1}{2}} \langle Y(\omega)Y^*(\omega) \rangle^{\frac{1}{2}}} \quad (4.28)$$

The information contained in this quantity is essentially an averaged phase difference for each frequency between the experimental signals $x(t)$ and $y(t)$. Clearly, because the phase difference is related to the angular spacing between coils and is assumed to be constant in the interval used to perform the Fourier transforms, whenever a coherent mode is present, $\chi(\omega)$ is essentially the centroid of a clusterized set of phase differences in the complex plane. On the other hand, truly uncorrelated signals at some frequency ω will exhibit a set of phase differences uniformly distributed over the complex plane in a way such that $\chi(\omega)$ is close to the origin in the same plane.

Coherence $\gamma(\omega)$, defined by (4.29), is the squared length of the array defined by the origin and $\chi(\omega)$ in the complex plane. It can have any value in the interval $[0 \ 1]$, where 1 corresponds to maximum correlation (no phase slippage) and 0 to complete absence of correlation. The average phase difference $\Delta\Psi$ between the two signals at frequency ω

is given by (4.30) where $\chi_i(\omega)$ and $\chi_r(\omega)$ are, respectively, the imaginary and real parts of $\chi(\omega)$. Finally, $\Delta\Psi(\omega)$ can be used in (4.26) to obtain the spatial periodicity.

$$\gamma(\omega) = \chi^2(\omega) \quad (4.29)$$

$$\Delta\Psi(\omega) = \arctan\left(\frac{\chi_i(\omega)}{\chi_r(\omega)}\right) \quad (4.30)$$

Usage of the popular FFT algorithm determines the frequency resolution to be $\Delta f = F_S/N$ where F_S is the sampling frequency and N is the number of samples used in the Fourier transforms. The temporal resolution of this method is configurable as (4.28) is applicable in the vicinity of any time t however, the temporal localisation is disperse throughout the interval used to perform the Fourier transforms.

A KF based implementation of the coherence spectrum for mode number analysis is reported in [Coel 08] where it is shown that, under real-time requirements and/or whenever a limited amount of samples is available for the analysis (for instance in ISTTOK due to its naturally bursty MHD activity), the proposed method is more suitable than the conventional Fourier approach.

Singular Value Decomposition Another approach often used for mode number analysis is the SVD [Tref 97]. Formally, the SVD is a matrix factorization of the form $X = USV^*$ where X is an $m \times n$ rectangular matrix, U is an $m \times m$ unitary matrix, S is an $m \times n$ diagonal matrix, V is an $n \times n$ unitary matrix and V^* denotes the conjugate transpose of V . S contains the (non-negative) singular values of X .

The SVD determines the matching set of basis vectors for each dimension of X in the directions of maximum coherence. Equivalently, it calculates the eigenvectors and eigenvalues of $\text{Cov}(X)$. One basis, the columns of U , is formed by the eigenvectors of X^*X and the other basis, the columns of V , is formed by the eigenvectors of XX^* . Their eigenvalues coincide and are equal to the square of X 's singular values.

Considering a set of n coils, located at the same poloidal position and different toroidal locations, for which a set of m synchronously sampled measurements is available, one can define the $m \times n$ rectangular matrix X as shown in (4.31).

$$X = \frac{1}{\sqrt{mn}} \begin{pmatrix} X_{1,1} & \cdots & X_{1,n} \\ \vdots & \ddots & \vdots \\ X_{m,1} & \cdots & X_{m,n} \end{pmatrix} \quad (4.31)$$

Factorizing X with the SVD method yields, separately, the temporal evolution and spatial structure of each coherent mode associated with particular singular values. The columns of U , also known as principal components, contain the temporal evolution and the columns of V , also known as principal axis, contain the spatial structure of the mode. Singular values provide the mode amplitude. More detailed discussions on the SVD method applied to mode number analysis in tokamak research are reported in [Nard 92, Kim 99, Raju 00].

Unlike the two Fourier-based methods presented earlier, SVD works in the time domain. Singular values are direct measures of the average mode amplitude within the data window however, further processing is required to extract both the frequency and spatial periodicity of the mode out of, respectively, matching columns of U and V . The DFT, for example, is a suitable candidate for this task.

One of the great strengths of the SVD method when applied to mode number analysis is that it is able to separate temporal from spatial periodicities of non-degenerate coherent structures thus simplifying the analysis. Still, frequency resolution depends on the method chosen to analyse the principal components. From a real-time implementation perspective, using a sliding SVD algorithm, temporal resolution is configurable and given by the temporal separation between the centers of two consecutive data sets used in the analysis however temporal localisation is, just like in the case of the coherence spectrum method, dispersed through the temporal data length.

Combined Fourier and SVD A combined Fourier and SVD method has been developed and applied to the analysis of coil data from the Mega Ampere Spherical Tokamak (MAST) [Hole 07]. To illustrate the proposed analysis let $X(\phi_i, t_k, \omega)$ be the time Fourier transform of coil i , located at toroidal angle ϕ_i , using measurements sampled in the interval $t \in [t_k - \delta t, t_k + \delta t]$. In addition, let one assume a spatial expansion of the previous in an orthogonal Fourier basis of the form $X(\phi_i, t_k, \omega) = \sum_n \alpha(n, t_k, \omega) e^{jn\phi_i}$ where $\alpha(n, t_k, \omega)$

denotes the complex-valued amplitude of toroidal periodicity n at angular frequency ω around $t = t_k$. Writing the previous in matrix form yields $X = \gamma \cdot \alpha$ with X , γ and α shown in (4.32) and (4.33). The matrix of complex mode amplitudes is therefore given by $\alpha = \gamma^{-1} X$ where the pseudo-inverse γ^{-1} is determined using SVD thus minimizing the residual $r = |\gamma \cdot \alpha - X|/|X|$.

$$X = \begin{pmatrix} X(\phi_1, t_k, \omega) \\ \vdots \\ X(\phi_N, t_k, \omega) \end{pmatrix}, \quad \gamma = \begin{pmatrix} e^{jn_1\phi_1} & \dots & e^{jn_M\phi_1} \\ \vdots & \ddots & \vdots \\ e^{jn_1\phi_N} & \dots & e^{jn_M\phi_N} \end{pmatrix} \quad (4.32)$$

$$\alpha = \begin{pmatrix} \alpha(n_1, t_k, \omega) \\ \vdots \\ \alpha(n_M, t_k, \omega) \end{pmatrix} \quad (4.33)$$

Subsequently, Monte-Carlo based simulations are used to study the statistical properties of the results assuming the coil measurements are normally distributed noise samples with mean and variance fitted to experimentally obtained data. This provides an estimation of the probability that the results obtained are reproducible by input noise. Time resolution and localisation and frequency resolution are determined by the sliding temporal Fourier transform of coil measurements as already discussed in section 4.4.3.

Periodogram based methods Least-square spectral analysis is a broad field of study which essentially comprises methods which perform, or are equivalent to, least-square fitting of sinusoidal sets of predefined frequencies to experimental data. These sort of methods are vastly used in cases where non-uniform data sampling occurs and standard Fourier analysis exhibits spectral aliasing artefacts. The Lomb periodogram [Lomb 76], demonstrated to be exactly equivalent to least-square fitting of sinusoids to experimental data [Scar 82], is probably the most popular of all in this family.

An interesting method based on the Lomb periodogram has recently been proposed for the analysis of Alfvén eigenmodes [Zege 06]. This method generalises the phase increment found in the usual formulation of Lomb's periodogram to also incorporate the spatial contribution leading to an expression for the spectral power as a function of the frequency and mode number(s). Furthermore, a time resolved version similar to what

is done in Wavelet analysis [Mall 99] is proposed to boost performance in the presence of non-stationarities. In particular, the size of the data window around a time instant is proportional to the inverse of the frequency being evaluated. This method requires a combinatorial evaluation of the spectral power in a frequency and mode number mesh. Time resolution is given by the temporal separation of the center of two consecutive data sets and time localisation is proportional to the inverse of the frequency being evaluated. No phase information can be extracted from this method.

Yet another recent approach used for toroidal mode number analysis of Alfvén eigenmodes in JET is based on the SparSpec method [S Bo 07]. Originally developed for the analysis of unevenly sampled astronomy and astrophysics experimental data, SparSpec is a linear least-squares method aiming to find the sparsest set of sinusoidal modes with non-zero amplitude that fit the experimental data. In practice, an approximate version of this intent is achieved by minimising the cost function $J(\mathbf{x}) = \frac{1}{2}\|\mathbf{y} - \mathbf{M}\mathbf{x}\|_2 + \chi\|\mathbf{x}\|_1$, a convex problem with no local minima, where $\|\cdot\|_p$ denotes the L^p -norm. In the case of toroidal mode number analysis, and for a particular time instant, \mathbf{y} is the set of M measurements for coils located at toroidal angles ϕ_j ($j = 1, 2, \dots, M$), \mathbf{x} is the array of N complex amplitudes to be estimated for each of the possible predefined set of modes $n_{min} \leq n_k \leq n_{max} \in \mathbb{Z}$ ($k = 1, 2, \dots, N$) and $\mathbf{M}_{jk} = e^{i2\pi\phi_j n_k}$. The parameter χ is a factor that essentially controls the weight given to sparsity when looking for the \mathbf{x} that minimizes J . A more detailed presentation of the method with application to experimental data can be found in [Klei 08a] and results of a real-time implementation of the algorithm in the same tokamak are reported in [Test 10].

Like the SDFT method discussed in section 4.4.3, SparSpec’s temporal resolution is the coil measurement’s sampling period and it does not provide any information on the mode frequency. It is a pure spatial decomposition. While for externally driven fluctuation scenarios, such as the resonant plasma response to the excitation of Toroidal Alfvén Eigenmodes (TAEs) by dedicated antennas, the mode frequency is externally imposed, for naturally occurring modes such as the NTM, the sawtooth precursor or fast particle driven TAEs there is no a priori information on the frequency of the mode for which the spatial (toroidal) periodicity is being inferred.

A new approach to real-time mode number analysis

Having summarily described some of the methods presently being used for mode number analysis, an innovative approach particularly suitable for a real-time implementation is now presented. This approach is based on the Kalman filter. Again, without loss of generality and taking into consideration the natural obstacles raised by poloidal harmonic coupling, the analysis presented herein is restricted to the estimation of the toroidal mode number.

A new KF based approach to real-time mode number analysis is summarized in figure 4.26. Magnetically induced voltages at the terminals of pick-up coil sensors are digitized by Analog-to-Digital Converters (ADCs) producing digital signals $S_k^{(m)}$ (where k denotes discrete time) for each of the $m = 1, 2, \dots, N_c$ coils located at toroidal positions ϕ_m . These signals are individually processed by non-linear KFs in the frequency tracking configuration (NLKFFT), see section 4.4.3. This process provides not only an estimate of the dominant tone’s frequency $\omega_k^{(m)}$ but also the in-phase component $x_{1,k}^{(m)}$ of $S_k^{(m)}$ in a finite frequency band around it, acting as a dynamic non-stationary band-pass filter. In addition a residue $r_k^{(m)}$, given by (4.34), is calculated for the past N samples providing a

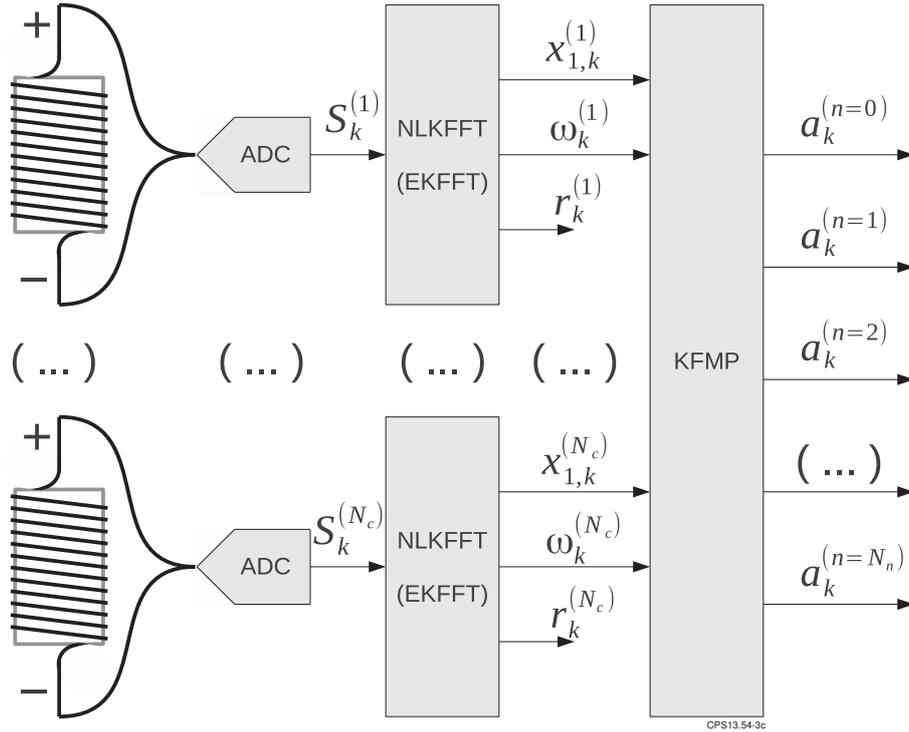


Figure 4.26: Data flow diagram for the KF based mode number analysis algorithm

crucial measurement of the estimation's accuracy.

$$r_k^{(m)} = \frac{\sum_{i=0}^{N-1} \left(S_{k-i}^{(m)} - x_{1,k-i}^{(m)} \right)^2}{\sum_{i=0}^{N-1} \left(S_{k-i}^{(m)} \right)^2} \quad (4.34)$$

Using the previously estimated $x_{1,k}^{(m)}$ as measurements, the KF Mode Projector (KFMP) implements a filter bank based on a predefined set of $N_n + 1$ toroidal mode numbers in order to determine the individual amplitudes $a_k^{(n)}$ of magnetic perturbations. Although not shown herein, information on the mode's propagating phase can be obtained with the same information that is used for the amplitude estimation. The frequency estimation $\omega_k^{(m)}$ performs a vital role in this linear time varying implementation of the KF as it provides the information that allows the modelling of the time evolution of the spatial periodicities, see section 4.4.3.

Non-linear Kalman Filter Frequency Tracker Non-linear Kalman filters for frequency tracking of magnetic signals in tokamak plasmas is the subject of [Alve 13a]. For the purpose of this work one can choose any of the non-linear KF methods presented therein. For simplicity, the EKFFT [La S 96] was chosen. Prior to [Alve 13a] the EKFFT had already been explored for fusion research applications in [Alve 10] and [Alve 11a]. The EKFFT employs model and measurement equations given, respectively, by (4.2) and (4.3), where z_k is the mean value of the measurement's probability density function (pdf) obtained by propagating the prior state pdf through the measurement equation. It is the likelihood in Bayesian terms, i.e., the most likely measurement given the prior state

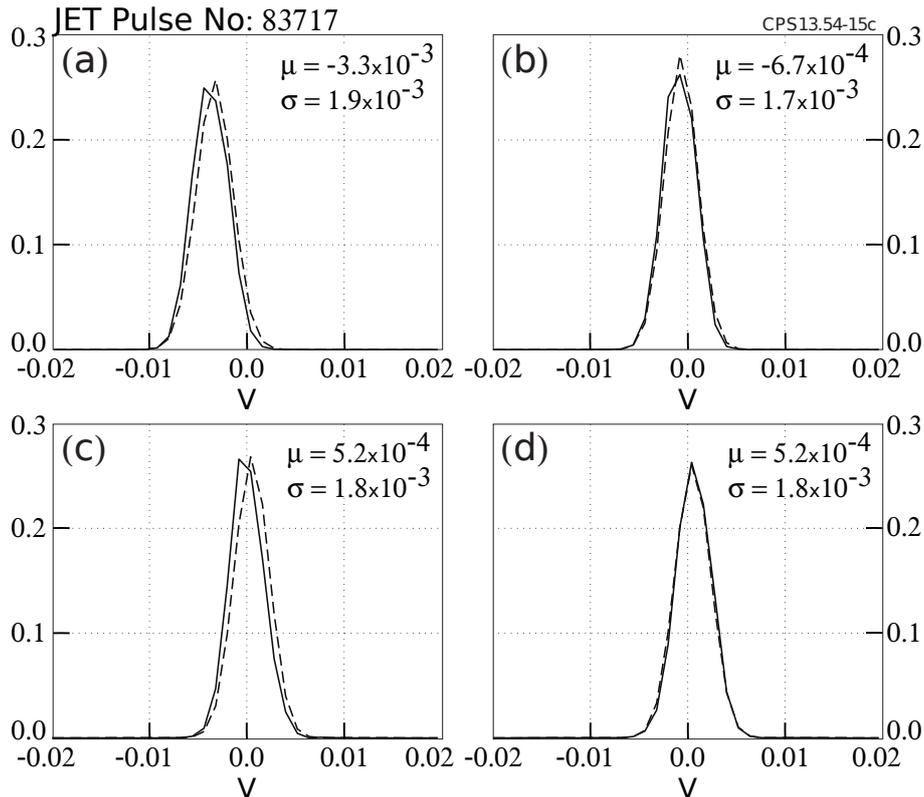


Figure 4.27: Noise pdf - experimental measurements (solid) and fitted gaussian (dashed) **a)** T002 coil located at $\phi = 42.15^\circ$; **b)** H305 coil located at $\phi = 110.38^\circ$; **c)** T007 coil located at $\phi = 222.17^\circ$; **d)** T009 coil located at $\phi = 290.37^\circ$.

pdf. In the EKF, all pdfs are approximated to normal distributions hence, the closer they truly are to normal distributions, the better the filter's performance.

Figure 4.27 shows, as a solid line, the measurement noise pdf determined experimentally for each of the four coils used throughout this work, see section 4.4.3 on the choice of the coil set. The data was obtained in the absence of plasma and energised circuits. The mean μ and standard deviation σ obtained from a normal pdf fit is also shown as a dashed line. The agreement between the measured and fitted parametrizations of a normal distribution is evident for each coil thus verifying, in these somewhat ideal conditions, the KF's assumption of white measurement noise. The coils are located at the same poloidal relative position ($\theta = 108.9^\circ$) and different toroidal locations ϕ .

The frequency, amplitude and phase of the dominant signal component can be obtained, respectively, from (4.4), (4.5) and (4.6) where F_s denotes the sampling frequency.

For the purpose of this work, $x_{1,k}$ and f_k are the required estimates. As demonstrated in [Alve 13a], more than one non-stationary frequency component can be isolated simultaneously. This feature allows, using the same number of parallel implementations of the KFMP applied to each of those isolated components, to retrieve information on the time evolution of the spatial periodicity of each. This type of implementation could be not only interesting while studying independent modes but also mode coupling processes. In this work the focus is solely on the dominant component.

Kalman Filter Mode Projector The idea behind this approach is to track the phase increment of each *band-passed* coil signal, monotonically in the counter clockwise toroidal

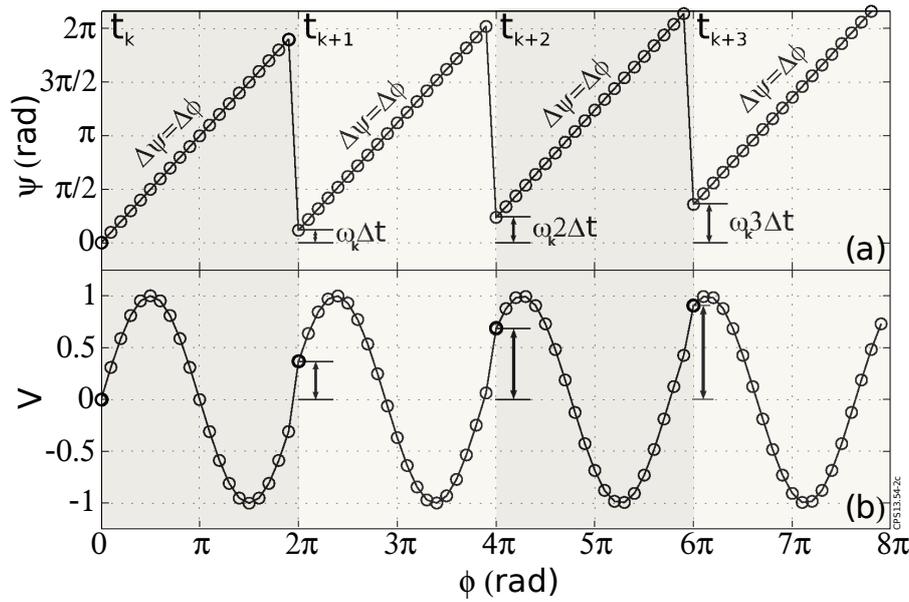


Figure 4.28: Spatio-temporal data unfolding mechanism - **a)** signal phase as a function of coil location for 4 consecutive instants; **b)** measured signal as a function of coil location for 4 consecutive instants.

direction, by projecting it onto a basis of predefined mode numbers. Starting from a particular coil, the tracking is performed spatially in the counter clockwise direction for $t = t_k$ until the coil immediately before the first one is reached. At this point the procedure is repeated but now for $t = t_{k+1}$. This process can be viewed as a spatial and temporal unwrapping of the experimental data and is depicted in figure 4.28. To obtain this figure a set of signals from 20 equally spaced coils was simulated to reflect the presence of a single rotating $n = 1$ mode of unitary amplitude oscillating at 15 kHz. Each coil signal was synchronously sampled at $t = t_k, t_{k+1}, t_{k+2}, t_{k+3}$ consistent with an acquisition rate of $F_s = 250$ kSamples/s. Starting from the first coil, located at $\phi = 0$, all coils are uniformly spaced with $\Delta\phi = 2\pi/20$. The top plot shows the phase Ψ of each coil's signal as a function of its toroidal location ϕ for four consecutive time instants. It is evident from this plot that $\Delta\Psi = \Delta\phi$ hence an $n = 1$ mode. In addition, because the mode is oscillating at a finite frequency, the phase difference in the signal of any fixed coil between two consecutive time instants is $\omega_k \Delta t$ where $\Delta t = 1/F_s = 4\mu s$. The bottom plot in the same figure shows the actual signal measured by each coil as a function of the coil's toroidal location for the same time instants. As expected, a single sinusoidal cycle for each time instant indicates the presence of an $n = 1$ spatial structure. Furthermore, the fact that this structure is being shifted to the left as time evolves, suggests that the mode is rotating. Equivalently, as indicated in the plot, this is merely the effect of the $\omega_k \Delta t$ phase jumps.

Let one now express this mechanism formally, in the context of KF estimation. Given a set of M coils located at toroidal angles ϕ_i , let their strictly positive angular spacing $\Delta\phi_i$ be defined by (4.35). Then, for a number of N predefined toroidal mode numbers $n = n_j$, the KF's non-stationary linear model is given by (4.36) and (4.37) and the state vector is $x = (x_{11} \ x_{12} \ x_{21} \ x_{22} \ \dots \ x_{N1} \ x_{N2})^T$ where x_{j1} and x_{j2} are respectively the in-phase and quadrature spatial components of a magnetic perturbation characterized by mode

number $n = n_j$ and angular frequency ω .

$$\Delta\phi_i = \begin{cases} \phi_i + (2\pi - \phi_M) & , \quad i = 1 \\ \phi_i - \phi_{i-1} & , \quad 1 < i \leq M \end{cases} \quad (4.35)$$

$$A_i = \begin{pmatrix} A'_{i1} & & \\ & \ddots & \\ & & A'_{iN} \end{pmatrix} \quad (4.36)$$

$$A'_{ij} = \begin{cases} \begin{pmatrix} \cos(\omega_i\Delta t - n_j\Delta\phi_i) & -\sin(\omega_i\Delta t - n_j\Delta\phi_i) \\ \sin(\omega_i\Delta t - n_j\Delta\phi_i) & \cos(\omega_i\Delta t - n_j\Delta\phi_i) \end{pmatrix} & , \quad i = 1 \\ \begin{pmatrix} \cos(-n_j\Delta\phi_i) & -\sin(-n_j\Delta\phi_i) \\ \sin(-n_j\Delta\phi_i) & \cos(-n_j\Delta\phi_i) \end{pmatrix} & , \quad 1 < i \leq M \end{cases} \quad (4.37)$$

If one now lets $i = ((k - 1) \bmod M) + 1$, where $k \in \mathbb{N}_{>0}$ is a counter of signal samples taken from each coil in the hybrid space and time sequence previously described and illustrated in figure 4.28, equations (2.31) to (2.35) together with (4.35)-(4.37) form the full set required for the analysis. Note that, in this case, k does not have the strict temporal significance commonly attributed to it in the KF formulation becoming, instead, a space-time counter or, equivalently, a cyclical sensor counter. The amplitude and phase, at index k , of a magnetic perturbation characterized by toroidal mode number $n = n_j$ is respectively given by (4.38) and (4.39). Finally, the measurement matrix \mathbf{C}_k relating the system state and the measurements, is the $1 \times N$ matrix obtained from concatenating the vector $(1 \ 0) N$ times, expressing the fact that a single measurement is the sum of all in-phase component estimations for each toroidal mode number considered.

$$a_{j,k} = \sqrt{x_{j1,k}^2 + x_{j2,k}^2} \quad (4.38)$$

$$\phi_{j,k} = \arctan\left(\frac{x_{j2,k}}{x_{j1,k}}\right) \quad (4.39)$$

As an illustration of one practical example of the KFMP, let one consider the case of $M = 6$ uniformly spaced coils in the toroidal direction whose measurements are synchronously sampled at the same rate of $F_s = 250$ kHz. Measurements $S_i(t, \omega(t))$ for each coil i were simulated to exhibit mode degeneracy, both $n = 1$ and $n = 2$ activity, at the same frequency $\omega(t)$ according to equations (4.40)-(4.42). $S_{n=x}(t)$ is the amplitude of the $n = x$ mode, $\Psi_{n=x}^i(t, \omega(t))$ is the total phase of the same perturbation as seen by coil i and $j = \sqrt{-1}$. The frequency $\omega(t)$, at which the perturbation is observed, is modulated at 20Hz in the interval 7.5-12.5 kHz as show in figure 4.29a. The amplitudes of the $n = 1$ and $n = 2$ modes are also modulated, respectively, at 40Hz and 53.34Hz in the interval $[0, 1]V$ as shown by the dashed lines in figure 4.29b. In addition, zero mean normally distributed noise with standard deviation $\sigma = 2 \times 10^{-3}V$ (compatible with coil measurements at JET, see figure 4.27) was added to each coil's signal. With this scenario, the KFMP was configured to simultaneously track $N = 4$ modes ($n = 0, 1, 2, 3$) with measurement variance $R = \sigma^2 = 4 \times 10^{-6}$ and $2N \times 2N$ process covariance $\mathbf{Q} = \chi\mathbf{I}$ where $\chi = 10^{-2}$ and \mathbf{I} is the identity matrix. The competence of the proposed method is demonstrated in figure 4.29b where the correct tracking of mode amplitudes is confirmed.

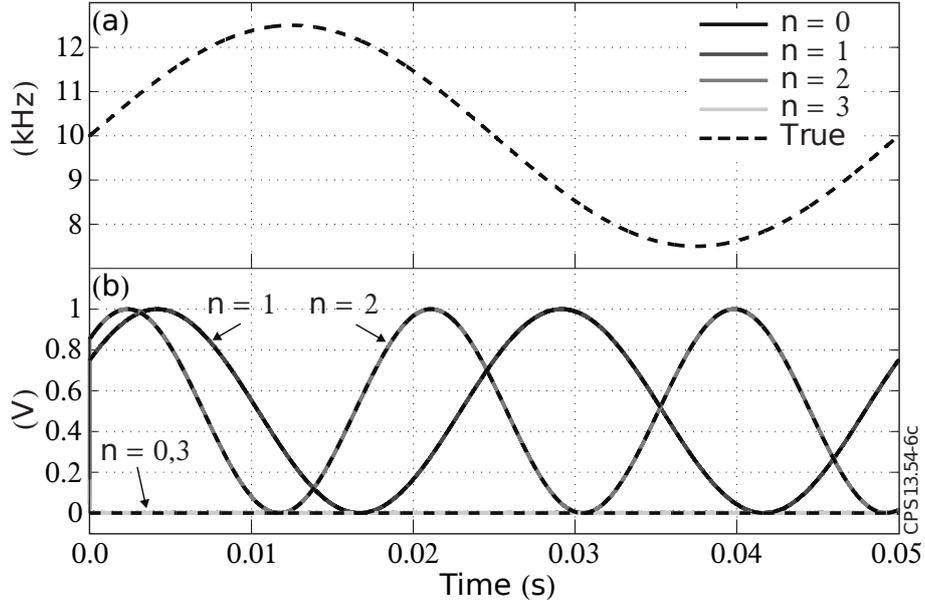


Figure 4.29: Degenerate modes of different modulated amplitudes observed at the same modulated frequency - **a)** Modulated frequency at which the modes are observed; **b)** Target (dashed) and estimated (solid) amplitudes for each mode.

Furthermore, in this case, amplitude estimations are provided at the input signals' sampling rate of $F_s = 250$ kHz.

$$S_i(t, \omega(t)) = S_{n=1}(t)e^{j\Psi_{n=1}^i(t, \omega(t))} + S_{n=2}(t)e^{j\Psi_{n=2}^i(t, \omega(t))} \quad (4.40)$$

$$\Psi_{n=1}^i(t, \omega(t)) = \int_0^t \omega(t') dt' - 1\phi_i \quad (4.41)$$

$$\Psi_{n=2}^i(t, \omega(t)) = \int_0^t \omega(t') dt' - 2\phi_i \quad (4.42)$$

In the previous example one has used $\chi = 10^{-2}$ which, along with $R = \sigma^2$ and an experimentally determined value of $\sigma = 2 \times 10^{-3}$ V for JET coils, was effective. For the same JET relevant conditions one can determine the χ dependence of the absolute value of the mean amplitude estimation error $|\bar{\epsilon}_j| = | \langle a_{j,k} - a_{j,k}^{true} \rangle |$, where $\langle \cdot \rangle$ denotes temporal averaging, for a magnetic perturbation characterized by a toroidal mode number $n = n_j$. This was achieved by repeating the KFMP estimation process of the previous example and calculating $|\bar{\epsilon}_j|$ for different values of $\chi \in [10^{-20} \ 10^{20}]$ in order of magnitude steps. A relevant zoom of these results is shown in figure 4.30. It can be seen that for $\chi \lesssim 10^{-15}$ the amplitude estimation error is large for all modes. This is because the filter trusts the model much more than the measurement hence, takes a long time to converge and is very slow to respond to measurement changes. In the case of the modes with non-zero amplitude ($n = 1$ and $n = 2$), for $10^{-14} \lesssim \chi \lesssim 10^{-7}$ it can be seen that the error gradually decreases as the filter starts trusting more in the measurement eventually saturating for $\chi \gtrsim 10^{-6}$. In the case of the modes with zero amplitude ($n = 0$ and $n = 3$), for $10^{-14} \lesssim \chi \lesssim 10^{-11}$ it can be seen that the error gradually decreases as the filter starts trusting more in the measurement thus converging faster to the correct amplitude. For $\chi = 10^{-11}$ there is a minimum because there is a good compromise between

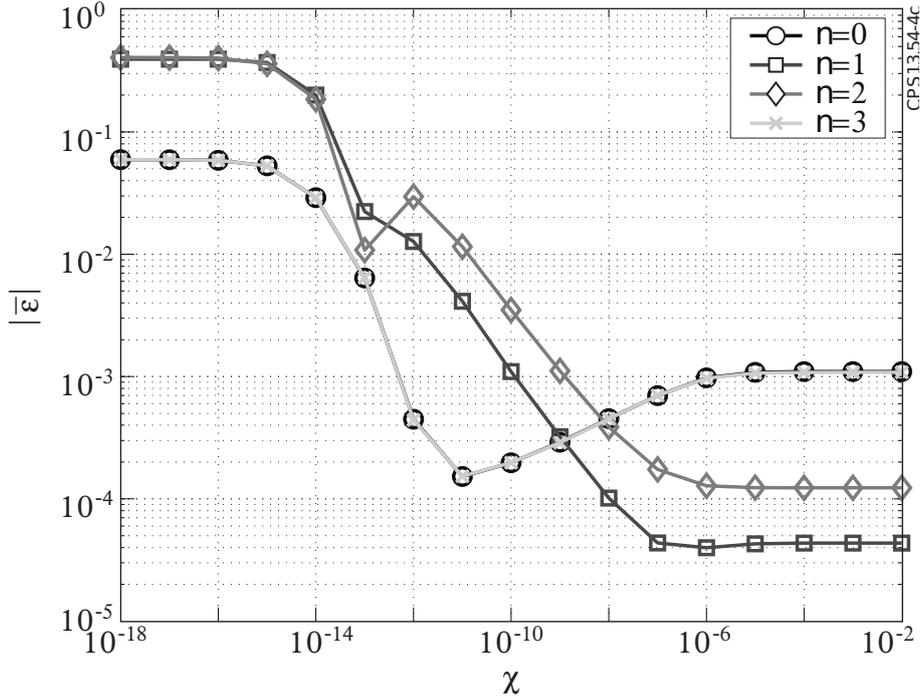


Figure 4.30: Absolute value of the mean amplitude estimation error as a function of the process covariance parameter χ .

noise rejection performance (better for small χ) and convergence time (better for large χ) to the correct stationary amplitude. Then, for $\chi \gtrsim 10^{-10}$, noise rejection deteriorates and, because the amplitude for these modes is stationary, a faster response does not help hence the error increases until it eventually saturates. Clearly, for JET conditions, a value of $\chi \in [10^{-6} \ 10^{-2}]$ is the best compromise between noise attenuation and non-stationary response in the filter's performance.

Sensor distribution at JET

Assuming a set of pickup coils located at the same poloidal position θ and at different toroidal angles ϕ_i , the perturbed magnetic field \tilde{B}_i at a particular time instant and pickup coil location is given by (4.43) where \tilde{B}_l , $m_l \in \mathbb{Z}$, $n_l \in \mathbb{Z}$ and Φ_l are respectively the amplitude, poloidal number, toroidal number and phase at zero poloidal and toroidal angles of mode l . This assumes the usual expansion in toroidal and poloidal harmonics of the magnetic perturbations measured at the plasma edge.

$$\tilde{B}_i = \sum_{l=1}^{\infty} \tilde{B}_l \exp [j(m_l \theta - n_l \phi_i + \Phi_l)] \quad (4.43)$$

Because all coils are assumed to be in the same poloidal position, the phase shift due to the poloidal topology of the mode is the same across all coil measurements. Therefore the phase difference, measured by two coils separated toroidally by $\Delta\phi$ radians, produced by a single mode with toroidal number n is $n\Delta\phi$. At JET, a set of eleven coils sharing the same poloidal position are toroidally located at $\phi_i = (2.97, 42.15, 77.00, 92.94, 103.11, 108.74, 110.38, 182.94, 222.17, 257.1, 290.37)$ degrees. Assuming that one is interested in tracking magnetic perturbations characterised by toroidal mode numbers $0 \leq n \leq 3$, including the $m/n = 1/1$ sawtooth precursor, the

core confinement deteriorating 3/2 NTM and the potentially disruptive 2/1 NTM, it is pertinent to establish the best sets of coils to use for this purpose. Also, for computational reasons, one is interested in processing, in real-time, the least number of coils whilst retaining the capability of reconstructing any mode of interest. For this purpose, several combinations of JET coils have been tested by simulating the presence of all $n = 0, 1, 2, 3$ modes with amplitudes of 1, 2, 3 and 4 respectively and observed at the same stationary frequency $f = 15$ kHz. Without loss of generality, the same propagation frequency is assumed for all modes. This degenerate case is the most challenging case in technical terms. The KFMP was configured with $R = 4 \times 10^{-6}$ and $\chi = 10^{-2}$, following the results of the previous section for non-stationary performance, and *instructed* to track all four mode numbers simultaneously. Simulations have verified that all k -combinations of 11 coils with $k < 4$ give rise to steady-state worst case amplitude estimation errors $> 10\%$. For the case of $k = 4$, totalling 330 combinations, the same error is depicted in figure 4.31a. It is clear from this plot that it is possible, for certain sets of coils, to have worst case amplitude estimation errors $< 0.1\%$. The reason why some coil combinations give good results while others don't is explained by the concept of observability [Kalm 59] which assesses the ability to reconstruct the initial state of a system given the measurement history. Formally, if the state of a linear system has dimension n , the $n \times n$ observability matrix \mathbf{O} can be defined. A system is then said to be observable if the row rank of \mathbf{O} is n , i.e., \mathbf{O} is full rank. In order to build \mathbf{O} one must take into account that the KFMP is a linear time-varying system, i.e., the process matrix \mathbf{A} is not the same for every observation cycle. In fact, \mathbf{A}_i is given by (4.36) and (4.37) where, in this case, $M = 4$ coils. Measurement vector $\mathbf{C} = (1 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1 \ 0)$ reflects the fact that the measurement for each cycle is the sum of all in-phase components of the four mode numbers being simultaneously considered. Hence, according to its standard definition but bearing in mind that in this case the linear process model is time-varying, \mathbf{O} is given by (4.44).

$$\mathbf{O} = \begin{pmatrix} \mathbf{C} \\ \mathbf{C}\mathbf{A}_1 \\ \mathbf{C}\mathbf{A}_1\mathbf{A}_2 \\ \mathbf{C}\mathbf{A}_1\mathbf{A}_2\mathbf{A}_3 \\ \mathbf{C}\mathbf{A}_1\mathbf{A}_2\mathbf{A}_3\mathbf{A}_4 \\ \mathbf{C}\mathbf{A}_1\mathbf{A}_2\mathbf{A}_3\mathbf{A}_4\mathbf{A}_1 \\ \mathbf{C}\mathbf{A}_1\mathbf{A}_2\mathbf{A}_3\mathbf{A}_4\mathbf{A}_1\mathbf{A}_2 \\ \mathbf{C}\mathbf{A}_1\mathbf{A}_2\mathbf{A}_3\mathbf{A}_4\mathbf{A}_1\mathbf{A}_2\mathbf{A}_3 \end{pmatrix} \quad (4.44)$$

Conveniently $|\mathbf{O}|$, the determinant of \mathbf{O} , can be used as a measure of the observability of the system. The closer $|\mathbf{O}|$ is from zero, the closer \mathbf{O} is from not being full rank and, consequently, the least observable the system is. $||\mathbf{O}||$, the absolute value of the determinant of \mathbf{O} , has been calculated for all possible 4-combinations of 11 coils and is shown in figure 4.31b. The correlation between both plots in figure 4.31 is evident. As $||\mathbf{O}||$ approaches zero, the worst case amplitude estimation error increases. Figure 4.31 is the basis for choosing a set of four coils with small worst case amplitude estimation error and, equivalently, large absolute value of the observability matrix determinant.

Results

This section presents the results obtained using the method described in section 4.4.3 for the analysis of experimental data from the JET tokamak. The results presented herein are

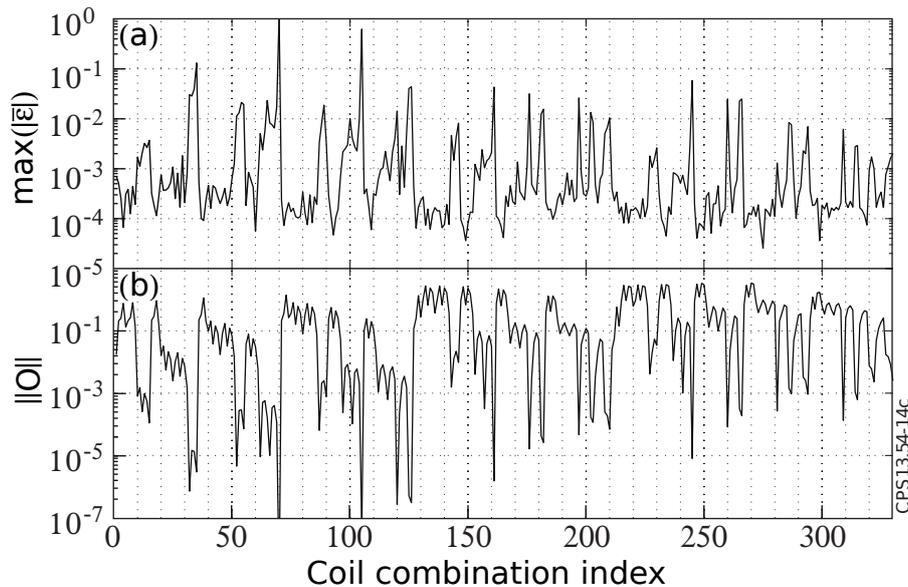


Figure 4.31: Observability of the system determines the estimation performance - **a)** worst case amplitude estimation error; **b)** absolute value of the observability matrix's determinant.

the result of the simultaneous analysis of the experimental data of four pick-up coils. One of the coil combinations with a worst case amplitude estimation error $< 10^{-3}$ and $\|\mathbf{O}\| > 1$ is a suitable choice. In this work, combination #132 corresponding to the coils located at toroidal angles $\phi_i = (42.15, 110.38, 222.17, 290.37)$ degrees was chosen. For the purpose of the work, a uniform sampling frequency of 250 kHz has been used decimating raw coil data where appropriate.

The EKFFT used herein was configured with $R = 4 \times 10^{-6}$ and the $3 \times 3 \mathbf{Q} = \mathbf{Q}_{diag} = (4 \times 10^{-8} \ 4 \times 10^{-8} \ 4 \times 10^{-10})$ matrix. Because, as demonstrated in [Bitt 00], the EKFFT gain depends solely on the ratio between the measurement and process covariances, the configuration of the EKFFT used in [Alve 13a] was updated to reflect the noise properties of the experimental data summarized in figure 4.27 whilst preserving the aforementioned ratio. Following the examples and studies presented in sections 4.4.3 and 4.4.3, the KFMP was configured with $R = 4 \times 10^{-6}$ and $\mathbf{Q} = \chi \mathbf{I}$ where $\chi = 10^{-2}$ and \mathbf{I} is the 8×8 identity matrix. Matrix dimensions are consistent with the simultaneous tracking of four toroidal mode numbers chosen to be $n = 0, 1, 2, 3$.

Figure 4.32 presents the results obtained for JET pulse 82635. This is a 1.3 MA, 1.3 T ELM-y H-mode pulse with 10 MW of additional power provided by Neutral Beam Injection (NBI). In figure 4.32a one can observe that the strongest MHD activity has an intermittent temporal evolution at frequencies below 10 kHz consistent with the typical observation from sawtooth instability precursors. The second harmonic of the sawtooth precursor can also be observed in the same plot. In figure 4.32b, mode number analysis based on the coherence spectrum is presented. A 2000 point (125 Hz frequency resolution) STDFT with 87.5% overlap is performed on the data of all four coils and the coherence and toroidal mode number spectrum is computed, as presented in section 4.4.3, with a temporal resolution of 20 ms. It is evident that the dominant harmonic of the mode exhibits an $n = 1$ toroidal topology. As expected, it can also be seen that its second harmonic is an $n = 2$. Figure 4.32c shows a comparison between the frequency estimation

given by the STDFT (black curve) and the average frequency estimation among the four EKFFTs (blue curve). As it should, these agree whenever the strongest frequency component's amplitude is significantly above the noise level. Figure 4.32d shows the KFMP's amplitude estimation for each of the $n = 0$ (blue curve), $n = 1$ (green curve), $n = 2$ (red curve) and $n = 3$ (cyan curve). The black curve shows, for the same non-stationary frequency tone, the amplitude estimated by the STDFT. It is evident, in accordance with the coherence analysis, that the dominant signal component is associated with a magnetic perturbation exhibiting an $n = 1$ toroidal topology. This observation is further substantiated by the fact that this is a well known perturbation typically characterised by its intermittent behaviour. Finally, figure 4.32e shows the residue of each of the EKFFT's estimations, given by (4.34). This quantity is crucial in terms of providing a measure of the estimation accuracy of the EKFFTs so that mode amplitudes subsequently calculated by the KFMP can be validated. It is clear that whenever the mode's amplitude is indistinguishable from the background noise level residues migrate to one and that otherwise the residues lie close to zero.

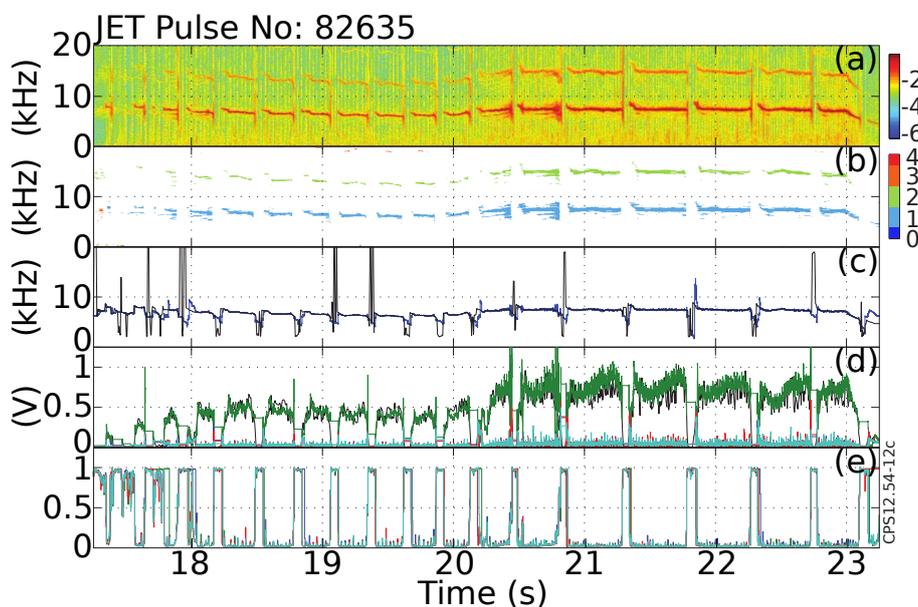


Figure 4.32: Toroidal mode number analysis - **a**) STDFT of a pick-up coil (\log_{10} scale); **b**) Mode analysis using the coherence spectrum; **c**) Frequency estimation given by the STDFT (black) and by the EKFFT (blue); **d**) Amplitude estimation given by the STDFT (black) and KFMP amplitude estimation of the $n=0$ (blue), $n=1$ (green), $n=2$ (red) and $n=3$ (cyan); **e**) EKFFT estimated residue for coil at toroidal position ϕ_1 (blue), ϕ_2 (green), ϕ_3 (red) and ϕ_4 (cyan).

Figure 4.33 shows the same type of analysis for JET pulse 79743. This is a 1.7 MA, 1.8 T H-mode pulse with 6 MW/1 MW of additional NBI/ICRH heating and low frequency strong type-I ELMs. This pulse is particularly interesting as it demonstrates the method's robustness in the presence of large frequency excursions of the order of ≈ 20 kHz in ≈ 1.5 s. In this case the dominant component is associated with a magnetic perturbation bearing a toroidal periodicity of $n = 2$ as corroborated by both the proposed method and the coherence spectrum approach. The very short-lived sawtooth activity, at ≈ 10 kHz and especially around 16-17.6 seconds and later around 20-23 seconds, has a small but often noticeable impact on the EKFFT's residues though never deteriorating the method's

performance on demand.

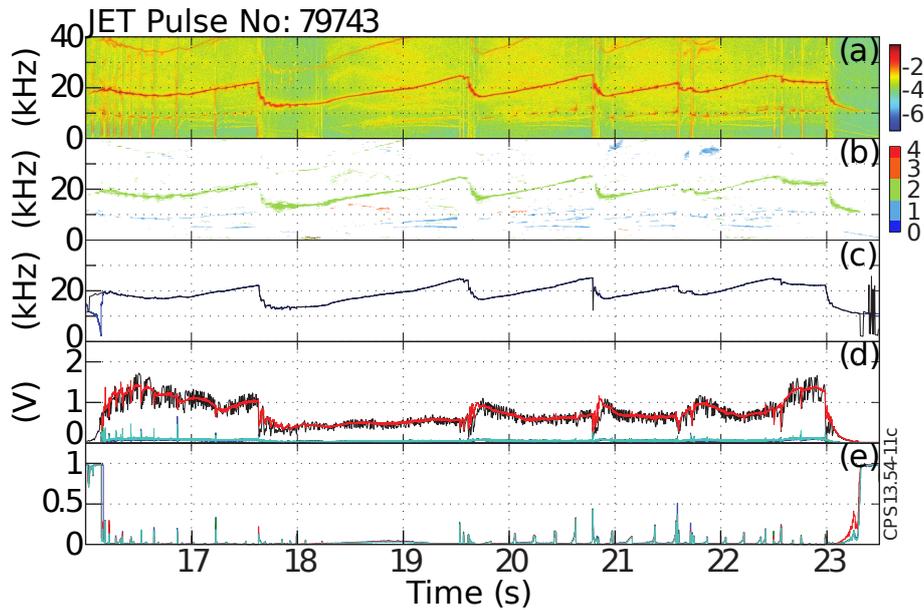


Figure 4.33: Toroidal mode number analysis - **a)** STDFT of a pick-up coil (\log_{10} scale); **b)** Mode analysis using the coherence spectrum; **c)** Frequency estimation given by the STDFT (black) and by the EKFFT (blue); **d)** Amplitude estimation given by the STDFT (black) and KFMP amplitude estimation of the $n=0$ (blue), $n=1$ (green), $n=2$ (red) and $n=3$ (cyan); **e)** EKFFT estimated residue for coil at toroidal position ϕ_1 (blue), ϕ_2 (green), ϕ_3 (red) and ϕ_4 (cyan).

As a final example, figure 4.34 shows the results obtained for JET pulse 83606. This is a 2 MA, 2.2 T ELM-y H-mode pulse with 11 MW of additional heating supplied by NBI. The dominant magnetic perturbation exhibits an $n = 2$ toroidal topology attested by both the KF based and coherence spectrum methods. This pulse demonstrates the resilience of the proposed method in the presence of frequent ELMs and both neighbouring and strong sawtooth activity. It is evident, however, the intermittent signature of ELMs in the residue and also appearing as a modulation in the estimated amplitude of the $n = 2$. Furthermore, the strong and moderately long-lived sawtooth activity at around 15.3-16.6 seconds, also evident in the residue, does not propagate to the amplitude estimation of the dominant component, as it should, since it is of a different toroidal topology.

This section has demonstrated the results obtainable with the proposed method for characteristic JET plasmas. Result validation has been performed by checking the agreement of this method with toroidal topologies inferred via the coherence spectrum method.

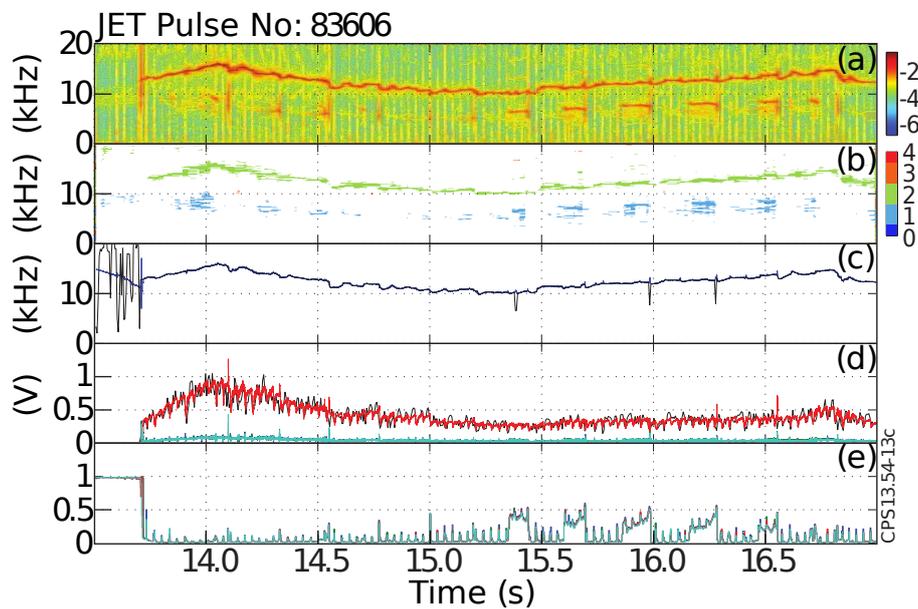


Figure 4.34: Toroidal mode number analysis - **a)** STDFT of a pick-up coil (\log_{10} scale); **b)** Mode analysis using the coherence spectrum; **c)** Frequency estimation given by the STDFT (black) and by the EKFFT (blue); **d)** Amplitude estimation given by the STDFT (black) and KFMP amplitude estimation of the $n=0$ (blue), $n=1$ (green), $n=2$ (red) and $n=3$ (cyan); **e)** EKFFT estimated residue for coil at toroidal position ϕ_1 (blue), ϕ_2 (green), ϕ_3 (red) and ϕ_4 (cyan).

Chapter 5

Control of external magnetic perturbations in the JET tokamak

External magnetic perturbations are typically utilized in tokamak devices with two operational or experimental purposes: the correction of the intrinsic 3D error fields due to possible misalignments of the toroidal field generating coils and the mitigation or suppression of ELMs [Lian 07]. Dedicated external coils (to the main plasma chamber) are used for the correction of the error field while for the mitigation of ELMs either external or internal coils can be used although the latter proves much more efficient. In fact, while for the correction of the error fields one is targeting the error fields existent even in the absence of plasma, in the mitigation of ELMs the additional driven magnetic fields are meant to ergodize slightly the magnetic topology in the pedestal region thereby reducing the drive for the destabilization of ELMs. This is facilitated if the driven magnetic fields are generated inside the main chamber, minimizing screening effects due to the conducting vessel and plasma rotation. The control of the magnetic field produced by these coils is obtained by controlling the current flowing in them. The real-time system responsible for this control underwent a number of improvements since its original implementation utilizing the present Voltage Controlled Voltage Source (VCVS) [Alve 11b].

This chapter describes the overall system, the built-in functionality and presents preliminary experimental results. In particular, the main improvements are: the possibility of automatically reduce the current references in case the plasma amplifies the perturbation applied [Grya 08], real-time limitation of dI/dt to limit electromotive force in machine protection systems, implementation of a Model Predictive Control (MPC) as an alternative to the Proportional Integral Derivative (PID) and the possibility of adapting the current reference, in real-time, using an external system.

5.1 Introduction

The use of non-axisymmetric magnetic perturbations is presently being explored in many tokamak devices such as DIII-D, ASDEX-U, MAST and JET. They serve, primarily, two main purposes: error field correction and ELM mitigation/suppression. The former tries to address and ameliorate the inevitability of magnetic field asymmetries due to coil imperfections and the presence of magnetic materials which can resonantly interact with plasma instabilities causing them to grow, lock (phase locking to the external static fields, also slowing down significantly the local plasma rotation) and ultimately cause the plasma to disrupt thus imparting significant heat loads onto facing components. The latter

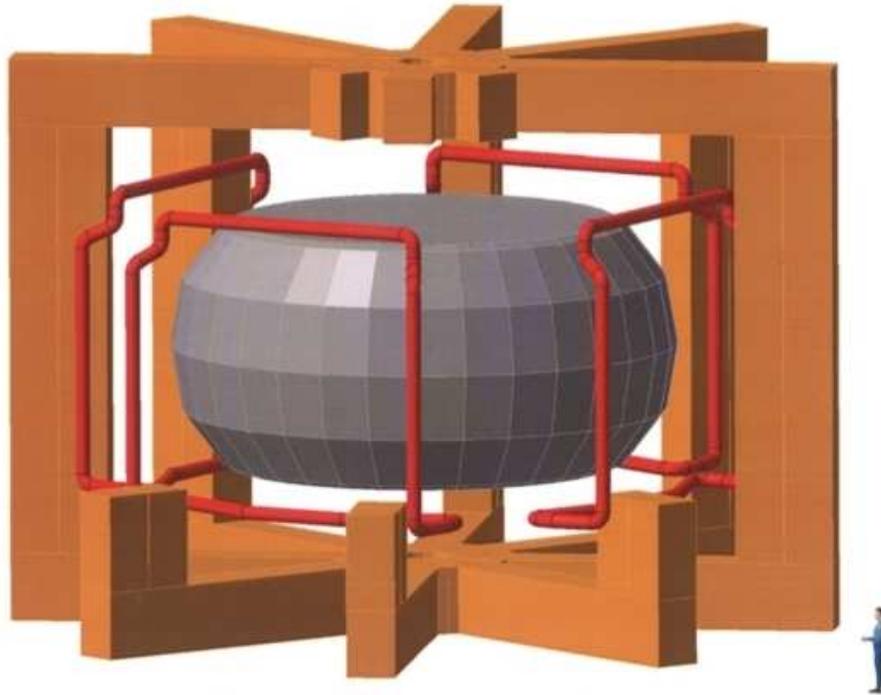


Figure 5.1: The error field correction coils (represented in red) at JET. *Courtesy of EFDA-JET.*

relies on the observation that, in certain operational regimes, applying external magnetic perturbations increases the frequency of ELM crashes (non-linear stage of the peeling-ballooning instability behind the ELM and driven by the plasma pedestal current density and pressure gradient) thus reducing their individual intensity and heat loads onto the facing components. Furthermore, low intensity time oscillating magnetic perturbations are also used for Resonant Field Amplification (RFA) studies diagnosing plasma stability near the so-called *no-wall limit* [Grya 08].

At JET, four coils placed 90° apart between the iron limbs and the plasma containing vessel are dedicated to *generating* external magnetic perturbations, see Fig. 5.1. Regardless of their present multi-purpose nature, they were historically baptised as the EFCCs. Current in these coils is driven by two dedicated VCVSs which, in turn, are controlled by the system presented herein.

5.1.1 Overview of the EFCC current controller system

The Error Field Correction Coils (EFCC) current controller system at JET has been implemented using Versa Module European (VME) technology. Further hardware details are given in section 5.2 and a basic data flow diagram of the system in JET's context is shown in Fig. 5.2. The system receives timing information, optically, from JET's Central Timing System (CTS) and uses JET's ATM based Real-Time Data Network (RTDN) [Felt 99] for a number of purposes: reporting control errors and operational limits violations to the Real-Time Protection Sequencer (RTPS) [Step 11], receiving current references from the Real-Time Central Controller (RTCC) [King 98] and measurements from magnetic diagnostics. The system also receives a series of analog measurements: plasma current and mode-lock signal from magnetics diagnostics as well as current and voltage measurements from the VCVSs. The analog outputs of the system are the VCVSs'

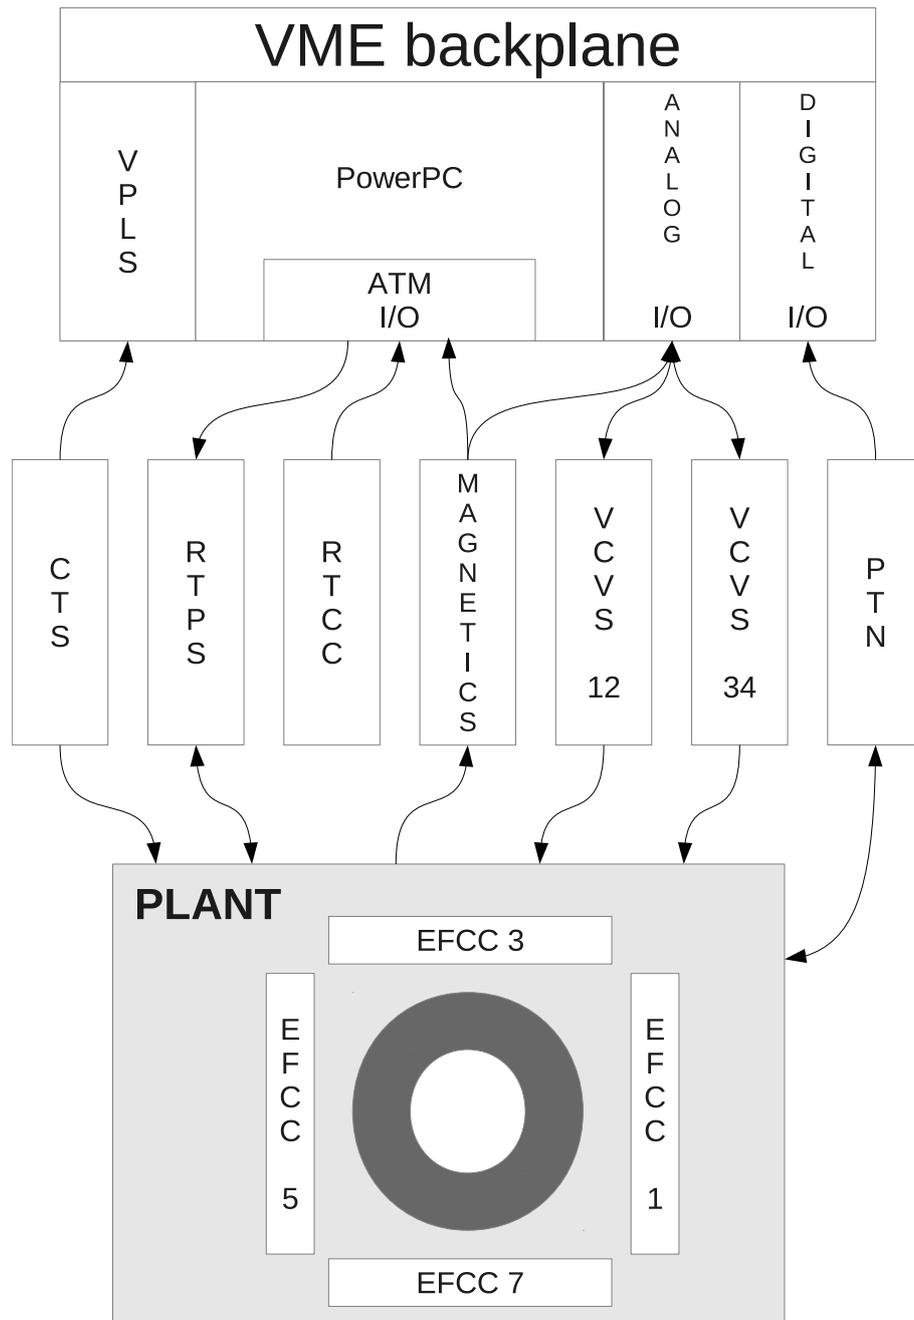


Figure 5.2: EFCC current controller in the JET context.

voltage requests. Finally, the system receives digital stop information from the Pulse Termination Network (PTN) [How 80].

5.1.2 User configuration

Typically at JET the person leading the experimental session, also known as the Session Leader (SL), pre-configures the target current waveforms for the EFCCs. In order to do this, a dedicated page is available in the ubiquitous pulse schedule editor [Beke 89] commonly known as *Level-1*. Using this time-window based graphical utility the SL is capable of designing the target current references (which can either be parametrized, e.g. sinusoidal, or point-based) configure protection related parameters, check the predicted I^2t

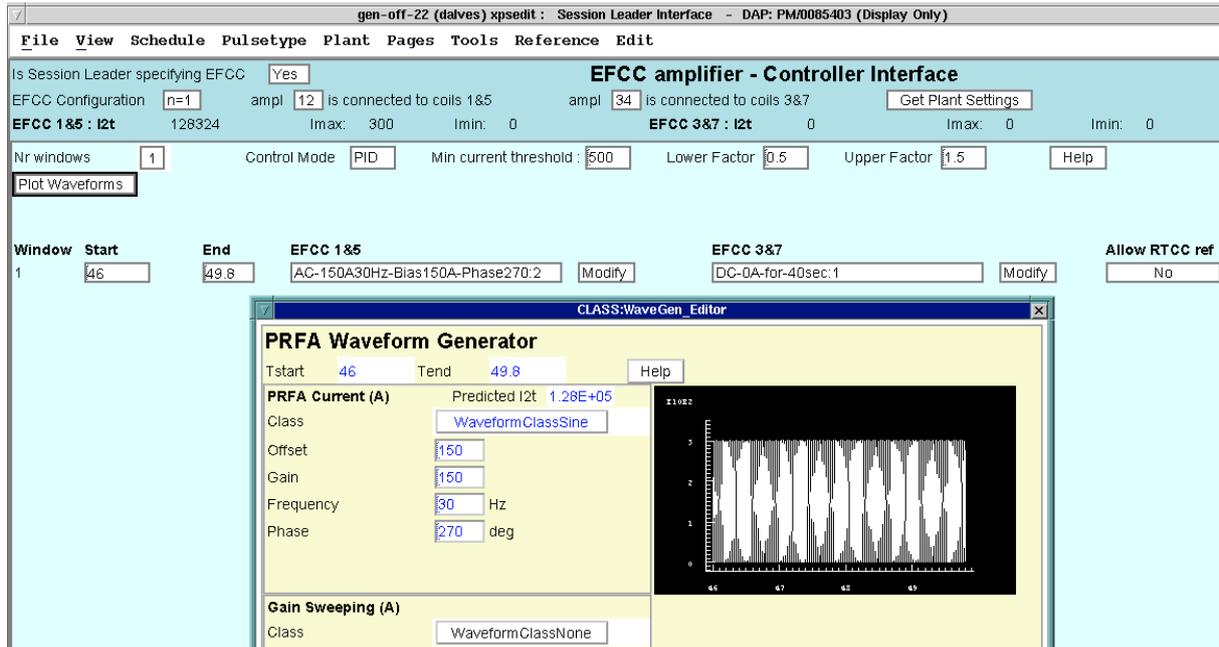


Figure 5.3: EFCC current controller Level-1 interface.

choose the control algorithm and decide for which time windows to enable control on references produced, in real-time, by RTCC.

5.2 Hardware

5.2.1 Error Field Correction Coils

As previously mentioned, the EFCCs [Barl 01] are 4 approximately square shaped coils (side length ≈ 6 m) located 90° apart around the exterior of the vessel, see Fig. 5.1. Each coil has 16 turns, a conductor cross section of 150mm^2 and is allowed to operate up to the I^2t limit of $252 \times 10^6 \text{A}^2\text{s}$. Some electromagnetic parameters of the EFCCs are presented in [Masi 02].

5.2.2 Voltage Controlled Voltage Sources

The VCVSs driving the EFCC current consist of two units, VCVS 12 and VCVS 34. The configuration of these 12-phase thyristor bridge rectifiers has evolved since they were first used at JET for the vertical stabilization of the plasma more than twenty years ago [Chir 91]. At that time they were baptised the poloidal radial field amplifiers (PRFAs). Presently they serve as two quadrant voltage controlled voltage amplifiers rated to provide -3 to $+3$ kV and 0 to 6 kA. Although the original manufacturer's documentation [Fina 82] specifies a bandwidth of 75 Hz, close to full voltage swings are limited by the 50 Hz mains supply either on the rising or falling edges depending on the polarity.

5.2.3 Controller system

The EFCC current controller is a VME based system containing the following cards:

VME Programmable Logic Service (VPLS)

JET specific card providing central timing and digital event information with 1 ms temporal resolution;

Motorola MVME5110

400 MHz PowerPC card with a total of 512 MB of RAM and on-board ethernet for slow monitoring, state transitions, pre-pulse configuration and post-pulse data collection;

Pentland MPV956

16 channel multiplexed ADC with a maximum acquisition rate of 250 kSample/s and 12 bit resolution ($[-10, 10]$ V) and 8 channel DAC with 12 bit resolution ($[0, 10]$ V);

Pentland MPV922

40 digital input channels and 32 digital output channels;

Interphase ATM NIC

ATM interface for real-time I/O.

5.3 Software

The combination of VME, PowerPC and the VxWorks Operating System (OS) has become one of the standards for real-time systems at JET. Therefore, adhering to this standard, the EFCC current controller system runs on top of version 5.5 of the VxWorks real-time OS.

5.3.1 Multithreaded Application Real-Time executor

An important spin-off of JET's distributed control philosophy is the Multithreaded Application Real-Time executor (MARTe) software framework for real-time systems [Neto 10]. Just like so many other products at JET, the MARTe framework is the result of nearly two decades of experience in designing and building software for real-time systems. MARTe is a C++ multi-platform framework for the development of modular and highly configurable real-time applications. At its core is the BaseLib2 library. This library implements a layered application programming interface (API) spanning basic low level (semaphores, mutexes, threads, filesystem interaction, high resolution timers, etc) to high level functionality such as messaging, http services and configuration parsing. It is very complete, optimised for real-time tasks and presently supports the Linux, Linux/RTAI, VxWorks, Solaris, Windows and the MacOSX OSs. Examples of the aforementioned diversity of MARTe based applications deployed at JET [De T 11] are the *vertical stabilisation* system (Linux/RTAI) [Bell 10], the previous version of the current controller of the EFCCs (VxWorks) [Alve 11b], the *real-time protection sequencer* (VxWorks) [Step 11], the *vessel thermal map* (Linux) [Alve 12a], the *wall load limiter system* (Linux), and the *hard X-ray and gamma-ray profile monitor* (Linux) [Fern 12]. The merits of the MARTe framework made it become, in the last few years, increasingly adopted and developed [Valc 12] by the magnetic confinement fusion community in europe [Neto 11]. It is also used in the ISTTOK [Carv 11], COMPASS [Valc 11] and FTU [Bonc 11] tokamaks and under consideration for ITER's fast plant system controllers [Gon 11].

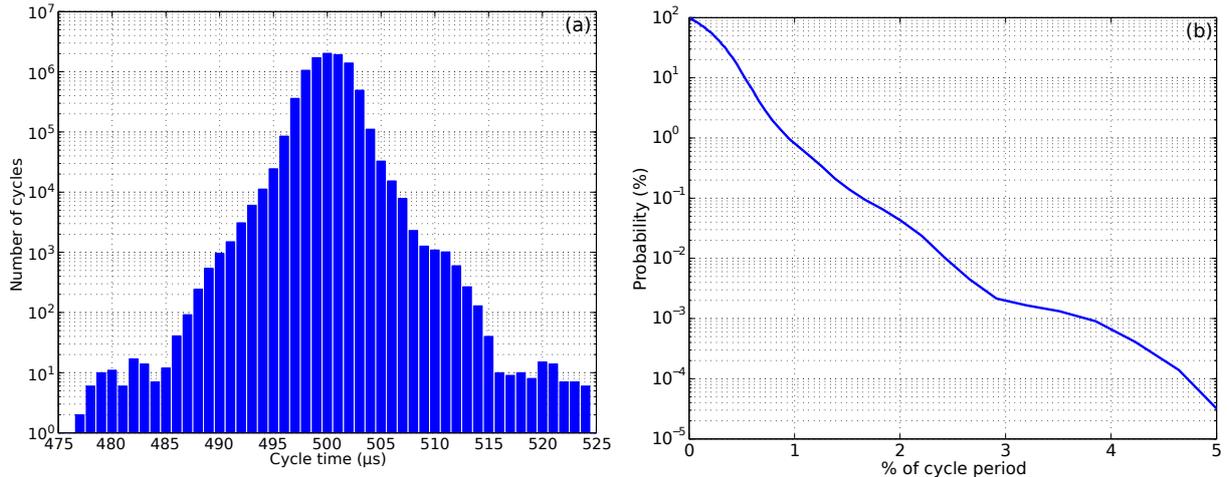


Figure 5.4: (a) Histogram of the application’s measured cycle time and (b) Probability of having a processing cycle with an absolute jitter larger than a given percentage of the application’s nominal cycle time.

One of its main strengths is that it allows application development, testing and debugging in non real-time environments and consequent online deployment without code changes because all the OS dependent functionality is abstracted and, therefore, transparent to the developer. Development, testing and commissioning time is minimised. Applications can also be developed in user-space (e.g. Linux) and deployed in kernel-space (e.g. Linux/RTAI or VxWorks) seamlessly. Being multi-threaded and particularly optimised for exploring multi-core architectures in a manageable and configurable way, dealing with the demands of modern real-time programming, e.g. using mutex priority inheritance and real-time schedulers, MARTe is not surprisingly the framework of choice for building high performance systems.

5.3.2 Synchronization

The processing time of each control cycle was optimized to run in under 500μs. Although JET’s central timing is made available to the system with a 1 ms resolution interrupting the CPU via a VME mailbox Interrupt ReQuest (IRQ), VxWorks’ auxiliary clock timer is able to provide similar functionality, within the CPU itself, at higher rates. Hence, the start of each real-time cycle is triggered by the *release* of a semaphore in the Interrupt Service Routine (ISR) handling the VxWorks timer IRQ configured to expire continuously at a rate of 2 kHz. The determinism of the application’s cycle time can be inferred from figure 5.4. In particular, from figure 5.4(b), it can be seen that the occurrence of a control cycle with an absolute jitter larger than 1% of the nominal application’s cycle time (500μs) is less than 1 every 100 cycles. In order to guarantee consistent time stamping of collected data with respect to JET’s central timing, whilst maintaining the 2 kHz control cycle, a scheme has been devised, see figure 5.5, making use of the PowerPC’s Time Stamp Counter (TSC), a 64 bit counter incremented every clock cycle. Hence, the sample time t_k attributed to signal samples either acquired or processed during the application’s k -th cycle is given by (5.1), where τ_k is the integer value of the TSC and χ_k is an estimate of the TSC’s period given by the exponential moving average $\chi_k = 0.9995\chi_{k-1} + 0.0005\chi_k^m$ and where $\chi_k^m = 10^{-3}/\Delta\tau^{(v)}$ is the latest available measurement of the TSC’s period. Also in (5.1), $t^{(v)}$ and $\tau^{(v)}$ are, respectively, the time given by the VPLS module (with a 1 ms

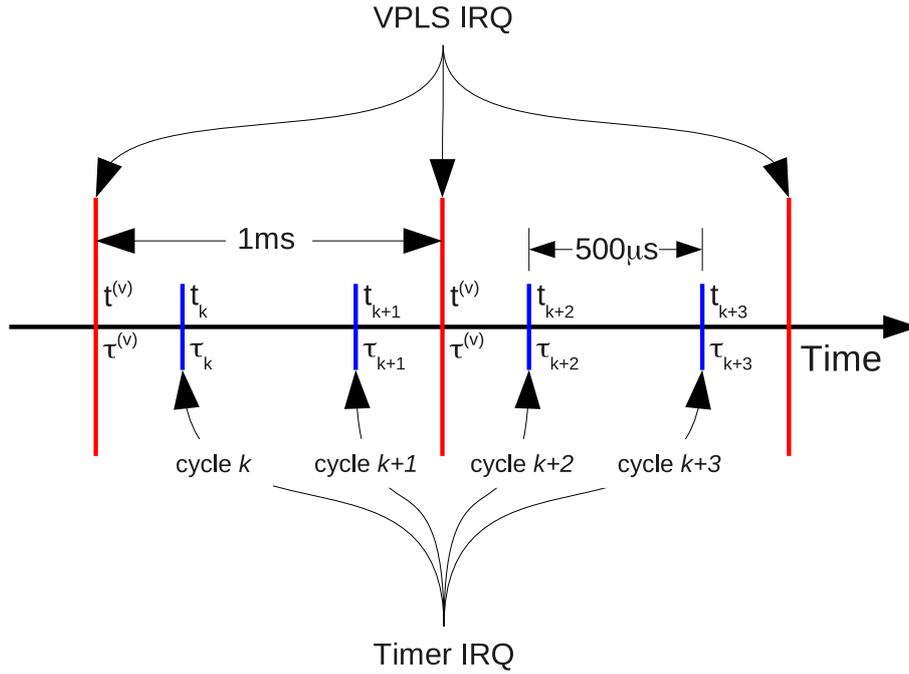


Figure 5.5: Time stamping and synchronisation mechanism.

resolution) and the value of the TSC when the VPLS IRQ immediately before the present control cycle occurred.

$$t_k = t^{(v)} + \chi_k (\tau_k - \tau^{(v)}) \quad (5.1)$$

5.3.3 Control modules

This section describes the various modules performing the global control task. A basic data flow diagram of the control algorithm is depicted in figure 5.6 henceforth used as reference. Detailed descriptions of the different modules and their interactions are presented below in individual subsections.

The plant itself essentially comprises the VCVS in an equivalent RL circuit containing the inductors (EFCCs) and a finite resistance. Using the Euler method to discretize a Thévenin equivalent RL circuit yields (5.2), where V_k is the output voltage of the VCVS, I_k is the current flowing in the EFCC circuit and R_{eq} , L_{eq} and $F_s = 2$ kHz are, respectively, the equivalent resistance, equivalent inductance and the sampling frequency of the system.

$$V_k = R_{eq}I_k + L_{eq}F_s (I_k - I_{k-1}) \quad (5.2)$$

Mutual inductances between the EFCC circuit and other circuits have been assessed and neglected in (5.2) to a good approximation. The mutual inductances were assessed by, while changing the currents in each of the other potentially coupled circuits individually, observing the voltage request necessary to control the current in the EFCC circuit to zero. It is worth noting that this was done using standard JET dry-runs and that there was no opportunity to perform systematic tests sweeping the rate of change of the currents in the various circuits thus inducing electromotive forces of different magnitudes in the EFCC circuit. Because, under these conditions, no mutual inductances have been measured, their effect was neglected in (5.2).

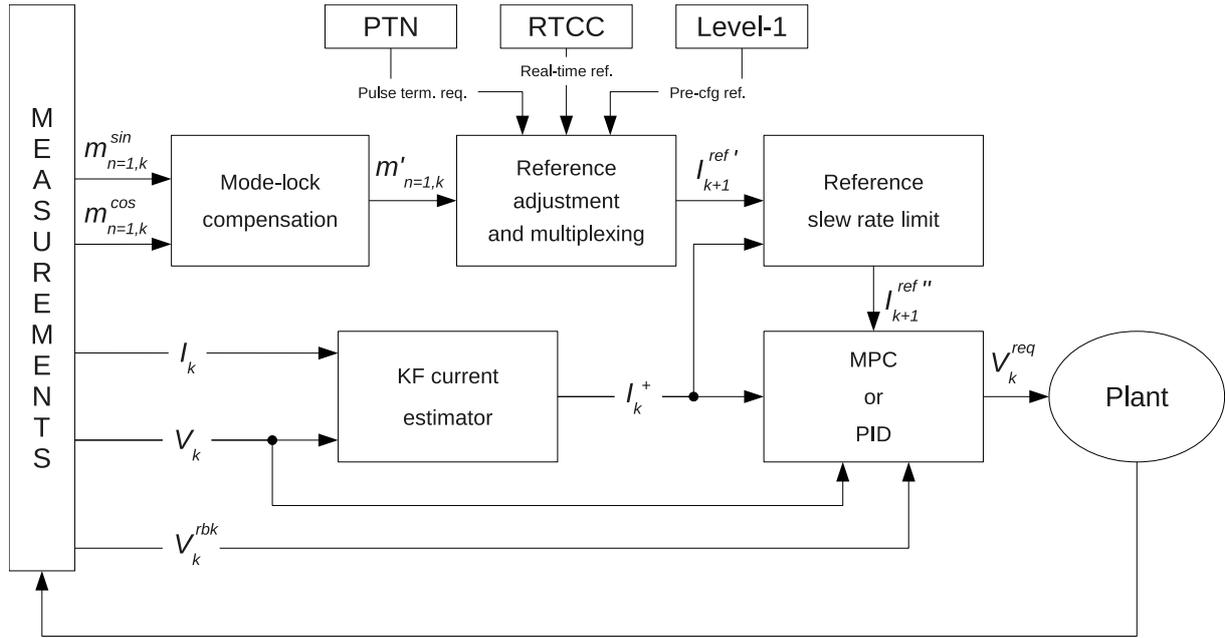


Figure 5.6: Control algorithm simplified data flow diagram.

All results presented in this dissertation were obtained with a single VCVS controlling the current in a circuit containing all four EFCCs. Basic analysis of circuit operation in steady-state conditions and of the exponential current decay $\propto \exp[-t/\tau]$, where $\tau = L_{eq}/R_{eq}$, establish that $R_{eq} \approx 330 \text{ m}\Omega$ and $L_{eq} \approx 36.7 \text{ mH}$.

Kalman filter current estimation

Using the circuit model (5.2), and given that experimental measurements of the EFCC current and VCVS output voltage are available, it is possible to use the KF to estimate the EFCC current with a better SNR than its actual measurement. In the light of (2.31) the prior EFCC current estimate I_k^- can therefore be written as (5.3) where I_{k-1}^+ is the posterior estimate of the previous cycle, V_k is a measurement of the VCVS's output voltage, $\alpha = L_{eq}F_s$ and $\beta = 1/(R_{eq} + \alpha)$.

$$I_k^- = \alpha\beta I_{k-1}^+ + \beta V_k \quad (5.3)$$

Similarly, in the light of (2.34), the posterior KF estimate of the EFCC current is given by (5.4).

$$I_k^+ = I_k^- + K_k \times (I_k^{meas} - I_k^-) \quad (5.4)$$

It is worth emphasizing that, as mentioned in the concluding remarks of section 2.2.1, the KF gain K_k needs not be computed in real-time since this is a linear time-invariant KF implementation. In order to pre-calculate the gain K_k one has used $R = 600$ (the measured noise variance in I_k^{meas}) and $Q = 10^{-1} \times R$ to obtain $K_k = 0.2673$ after recursively evolving (2.32), (2.33) and (2.35) until a steady-state is reached where K_k is considered to be sufficiently stationary. Figure 5.7 shows a comparative example between the raw current measurement I_k^{meas} and the KF's posterior current estimate I_k^+ where the SNR improvement is evident. The ratio between the variance of the noise in the KF's

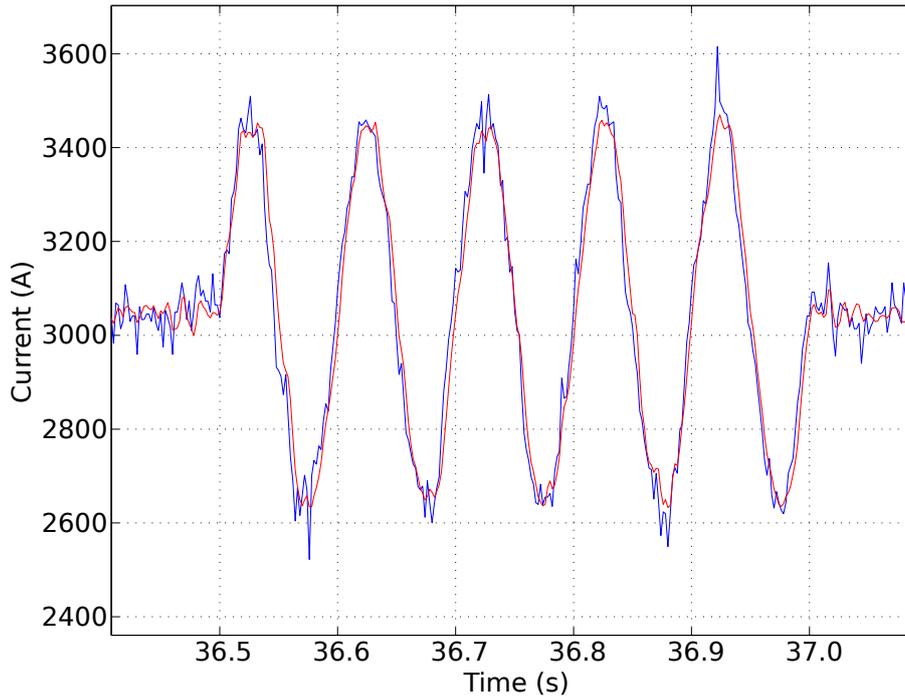


Figure 5.7: Time trace - (blue) Current measurement and (red) KF current estimate.

posterior current estimate and the variance of the noise in the current's raw measurement has been found to be $\lesssim 1/6$.

In order to further assess the performance of the KF current estimate, the frequency response in terms of amplitude and phase delay, when compared to the raw measurements, has been evaluated. Figure 5.8 shows the results of this study for a frequency scan of 10 – 70 Hz in steps of 10 Hz. The DFT is used to calculate the amplitude of both the KF estimate and the raw measurement. Figure 5.8(a) shows the difference between the two, normalized to the amplitude of the raw measurement. It can be seen that, up to 70 Hz, the amplitude attenuation of the KF's estimate never exceeds 10% of the actual amplitude. Figure 5.8(b) shows the delay between the KF's estimate and the current measurement, normalized to the period of the oscillation. The results were obtained using standard cross-correlation analysis of both signals and it can be seen that the delay of the KF's estimates never exceeds 7% of the oscillation's period.

Using this scheme and I_k^+ rather than I_k^{meas} for feedback control one reduces the propagation of the measurement noise into the control action itself. The cost to pay is some amplitude and phase distortion of the current estimate with respect to the measurement.

Mode-lock compensation

The mode-lock signal is commonly used for protection in tokamak research. It measures the amplitude of the potentially disruptive $n = 1$ MHD activity. At JET this signal is given by (5.6) and its components are given by (5.5), where k denotes the discrete time index. Each $S_{x,k}^{Octy}$ signal is the magnetic field measurement at time index k obtained by

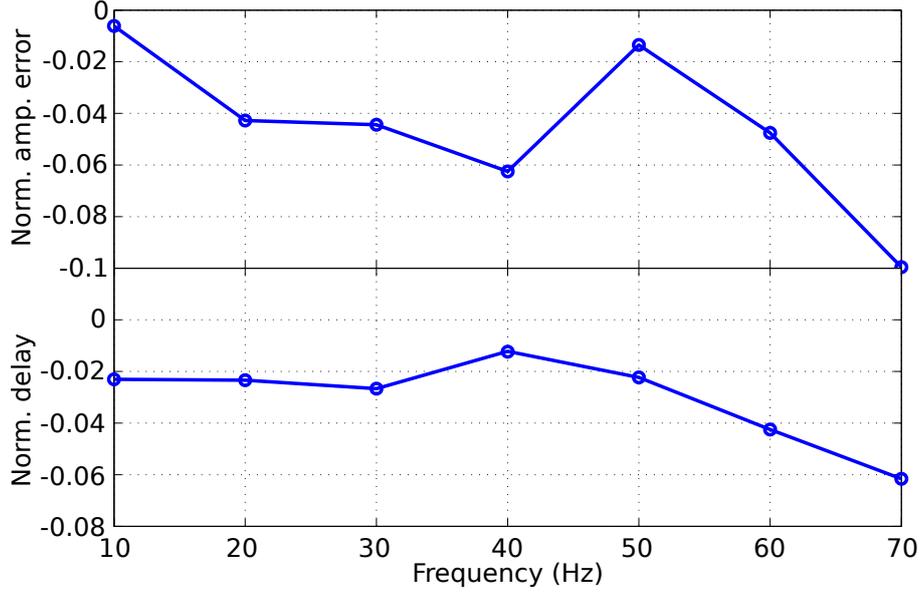


Figure 5.8: KF frequency response - (a) Normalized amplitude error and (b) Delay normalized to the period.

integrating saddle loop x located in octant y , see figure 5.9.

$$m_{n=1,k}^{\sin} = [S_{1,k}^{\text{Oct}_3} + S_{14,k}^{\text{Oct}_3}] - [S_{1,k}^{\text{Oct}_7} + S_{14,k}^{\text{Oct}_7}] \quad (5.5)$$

$$m_{n=1,k}^{\cos} = [S_{1,k}^{\text{Oct}_1} + S_{14,k}^{\text{Oct}_1}] - [S_{1,k}^{\text{Oct}_5} + S_{14,k}^{\text{Oct}_5}]$$

$$m_{n=1,k} = \sqrt{[m_{n=1,k}^{\sin}]^2 + [m_{n=1,k}^{\cos}]^2} \quad (5.6)$$

It is often the case that the plasma amplifies these RMPs to the point where the mode eventually locks to the resistive wall and disrupts the plasma imparting large eddy currents and forces on the containment vessel. One way to ameliorate this effect would be to, as soon as the $m_{n=1,k}$ is observed to reach dangerous values, automatically lower the current in the EFCCs to try and reduce the plasma amplification. However, in order to do that, it is assumed that this signal measures exclusively the plasma response and not the contribution of the direct pick-up from the external magnetic perturbation. It is therefore clear that some means to compensate $m_{n=1,k}$, so that it contains exclusively the contribution of the plasma response, is necessary. The form for the compensated mode-lock signal $m'_{n=1,k}$ used herein is shown in (5.8) and its, individually compensated, components are given by (5.7).

$$m_{n=1,k}^{\sin'} = m_{n=1,k}^{\sin} - \left(\sum_{i=0}^{N-1} a_i I_{k-i}^{\text{EFCC}_3} - \sum_{i=0}^{N-1} b_i I_{k-i}^{\text{EFCC}_7} \right) \quad (5.7)$$

$$m_{n=1,k}^{\cos'} = m_{n=1,k}^{\cos} - \left(\sum_{i=0}^{N-1} c_i I_{k-i}^{\text{EFCC}_1} - \sum_{i=0}^{N-1} d_i I_{k-i}^{\text{EFCC}_5} \right)$$

$$m'_{n=1,k} = \sqrt{[m_{n=1,k}^{\sin'}]^2 + [m_{n=1,k}^{\cos'}]^2} \quad (5.8)$$

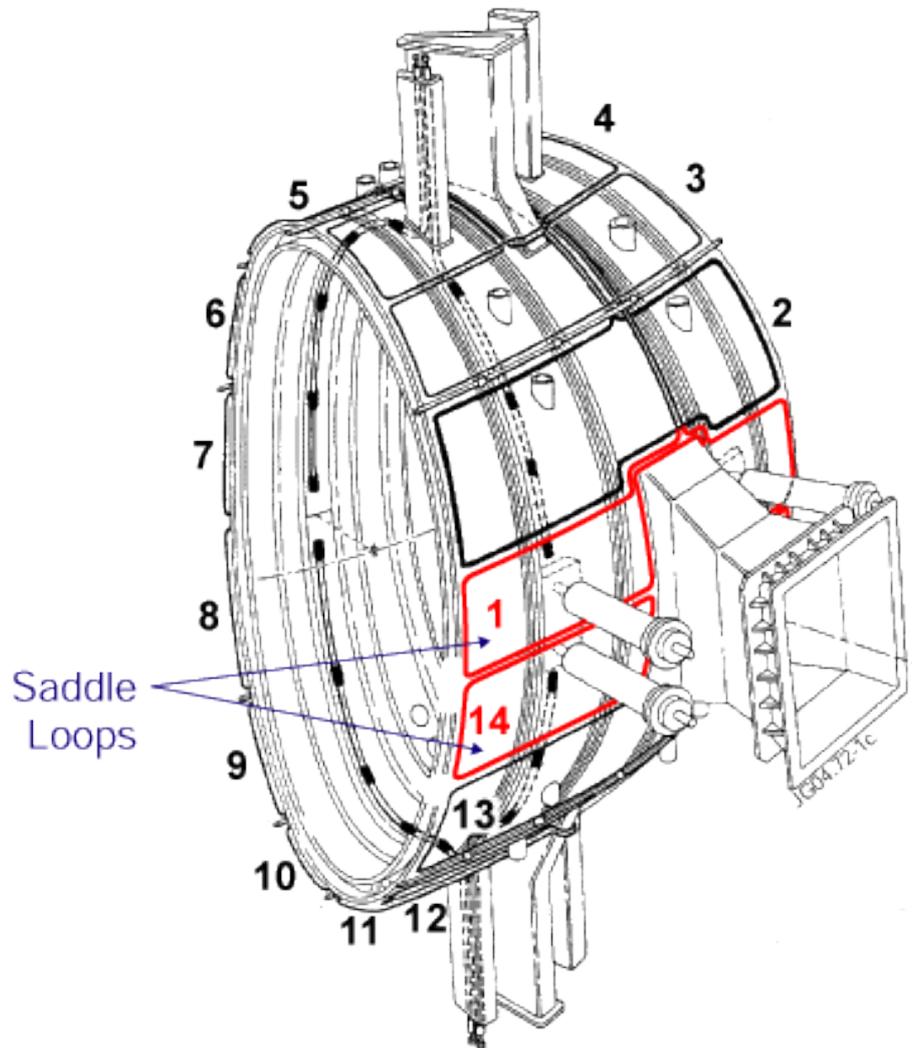


Figure 5.9: Saddle coils used in each of the 4 equally spaced octants to calculate the mode-lock signal. *Courtesy of EFDA-JET.*

The N coefficients in (5.7) were obtained by fitting (linear least-squares), in a discharge without plasma, $m_{n=1,k}^{sin'}$ and $m_{n=1,k}^{cos'}$ to zero for a 4 s window (72 – 76 s as shown in figure 5.10), during which no other circuits were active. The continuous lines represent the uncompensated signals ($m_{n=1,k}^{sin}$ and $m_{n=1,k}^{cos}$) while the dots represent the compensated ones ($m_{n=1,k}^{sin'}$ and $m_{n=1,k}^{cos'}$). For computational reasons the number N of coefficients in (5.7) should be as small as possible. However, using $N = 1$ implies completely disregarding any eddy currents induced, e.g. in the vessel, due to changes in the EFCC currents. In fact it can also be seen in figure 5.10 that, for the chosen value of $N = 2$, a better compensation is achieved in the region where the currents are stationary, i.e. $m_{n=1,k}^{sin}$ and $m_{n=1,k}^{cos}$ are essentially constant. It is worth stressing that this is not meant to be an exact cancellation mechanism but more of an approximate amelioration scheme.

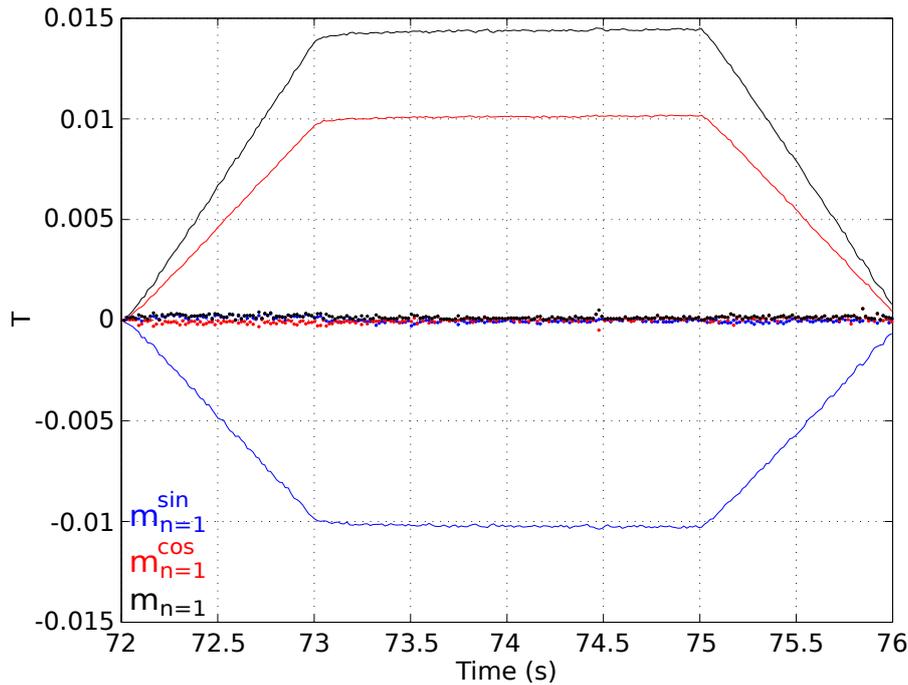


Figure 5.10: Mode-lock compensation (color legend in the figure) - continuous line represents the raw signals and dots represents the compensated signals.

Reference multiplexing, adaptation and checks

This module is responsible for: (a) checking that the EFCC currents are within pre-determined boundaries; (b) performing the multiplexing¹ of the references pre-configured by the SL and the real-time references from RTCC and (c) adjusting the references according to events in PTN and the compensated mode-lock signal calculated in the previous module.

After multiplexing, the current references I_{k+1}^{ref} are set to zero in the event of a plasma termination request or are adjusted according to (5.10), where γ_k is given by (5.9). m_0 and Δm are configurable parameters and their effect on γ_k is exemplified in figure 5.11.

$$\gamma_k = \frac{1}{2} \left[1 + \tanh \left(\frac{4(m_0 - m'_k)}{\Delta m} \right) \right] \quad (5.9)$$

$$I_{k+1}^{ref'} = \gamma_k \times I_{k+1}^{ref} \quad (5.10)$$

This mechanism significantly reduces the reference currents while an $n = 1$ mode of considerable amplitude exists in the plasma. As soon as the mode amplitude is reduced, the references automatically return to the desired values.

Reference slope limitation

This module is responsible for implementing a maximum absolute rate of change (slew rate) in the EFCC currents. The reason for this is because the toroidally asymmetric fields are inductively picked-up by the protection system of the toroidal field coils circuit (the

¹In the context of the present chapter, the term *multiplexing* is used to describe the selection process by which the current reference used in the control algorithm is chosen between the one pre-configured by the SL and the one made available by RTCC in real-time.

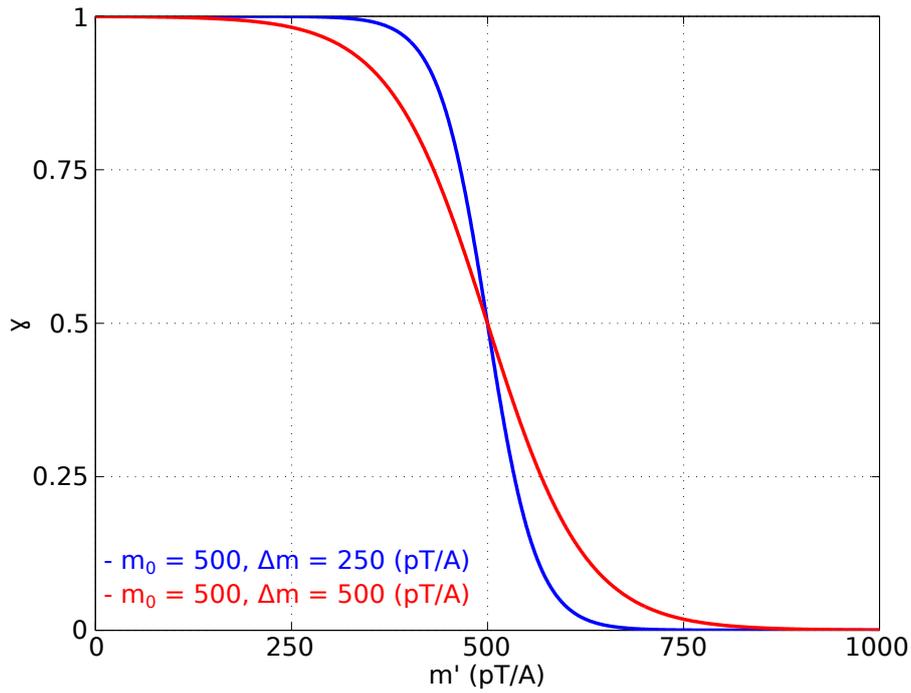


Figure 5.11: Automatic current reference adjustment based on mode-lock compensation.

direct magnet safety system) leading it to believe that a non-existing problem has occurred. It has been determined experimentally that the pick-up threshold for triggering an alarm is when the EFCC currents are varying at a rate $dI/dt \gtrsim 40$ kA/s. The idea is therefore to calculate the slope obtained through a linear regression of the past N current estimates plus the current reference $I_{k+1}^{ref'}$ desired to be achieved in the next control cycle¹. Hence a linear regression of $N + 1$ points. If the calculated slope a , given by (5.11)², is such that $|a| < |a_{th}|$, where $|a_{th}| = 40$ kA/s, then the new slope limited reference is the same as the adjusted reference calculated in the previous module, i.e. $I_{k+1}^{ref''} = I_{k+1}^{ref'}$. In case this condition is not satisfied, then it is possible to *reverse* (5.11) and calculate the new reference $I_{k+1}^{ref''}$ that satisfies $|a| = |a_{th}|$, thus optimizing the ramp-down rate to the fastest possible value without violating the slope threshold previously established. In this case $I_{k+1}^{ref''}$ is given by (5.12) where a_{th} is either ± 40 kA/s depending on the slope violation

¹Given two sets of data points $\{x_i, y_i\}$, with $i = 1, 2, \dots, M$, the slope of the straight line $y = ax + b$ that best fits $\{x_i, y_i\}$ in the least-squares sense, i.e. $\min_{a,b} \sum_{i=1}^M (y_i - ax_i - b)^2$, is given by $a = \text{Cov}[x, y] / \text{Var}[x]$.

² I_{k-i}^+ is the history of current estimates calculated by the KF module using (5.3) and (5.4) and $\Delta t = 1/F_s = 1/2000 = 500 \mu\text{s}$.

direction.

$$a = \left\{ N\Delta t \times I_{k+1}^{ref'} + \sum_{i=0}^{N-1} i\Delta t \times I_{k-i}^+ - \frac{1}{N+1} \left(\sum_{i=0}^N i\Delta t \right) \times \left(I_{k+1}^{ref'} + \sum_{i=0}^{N-1} I_{k-i}^+ \right) \right\} \times \quad (5.11)$$

$$\left\{ \sum_{i=0}^N (i\Delta t)^2 - \frac{1}{N+1} \left(\sum_{i=0}^N i\Delta t \right)^2 \right\}^{-1}$$

$$I_{k+1}^{ref''} = \left\{ a_{th} \times \left[\sum_{i=0}^N (i\Delta t)^2 - \frac{1}{N+1} \left(\sum_{i=0}^N i\Delta t \right)^2 \right] + \frac{1}{N+1} \left(\sum_{i=0}^N i\Delta t \right) \times \left(I_{k+1}^{ref'} + \sum_{i=0}^{N-1} I_{k-i}^+ \right) - \sum_{i=0}^{N-1} i\Delta t \times I_{k-i}^+ \right\} \times \left\{ N\Delta t - \frac{1}{N+1} \sum_{i=0}^N i\Delta t \right\}^{-1} \quad (5.12)$$

An example of this mechanism in action can be seen in figure 5.12 for a pulse where there was a termination request at $t \sim 73.55$ s. On every cycle, the last $N = 5$ KF current estimates plus the current reference for the next cycle ($I_{k+1}^{ref'}$) are used in linear regression's slope calculation. It can be seen from figure 5.12 that at $t \sim 73.55$ s the adjusted reference $I_{k+1}^{ref'}$ (in green) is immediately set to zero however, the actual reference used for control $I_{k+1}^{ref''}$ (the slope limited reference shown in red) exhibits a descent slope ~ 40 kA/s as desired. The latter is subsequently used to perform the desired controlled current ramp-down as depicted by the blue curve.

A much simpler mechanism could have been devised to perform the same task however, a significant advantage of the proposed method is that it is suitable for coping with the noisy current estimates. So, instead of using two noisy estimates to calculate a slope, using this method, and assuming that the KF current estimates are sufficiently unbiased (i.e. the noise associated with them is zero-mean), ensures that the slope calculation is obtained using the straight line that fits best the last N estimates plus $I_{k+1}^{ref'}$.

It is also worth noting that, if this mechanism is to be used not only when there has been a pulse termination request, the choice of N may impact on the overall bandwidth of the system. For example, in the case just shown where $N = 5$, means that $\Delta T = N/F_s = 5/2000 = 2.5$ ms which corresponds to a limiting frequency $f_l = 1/\Delta T = 400$ Hz, well above the overall system bandwidth.

5.3.4 Control modes

The user can choose between a PID and a MPC based algorithm.

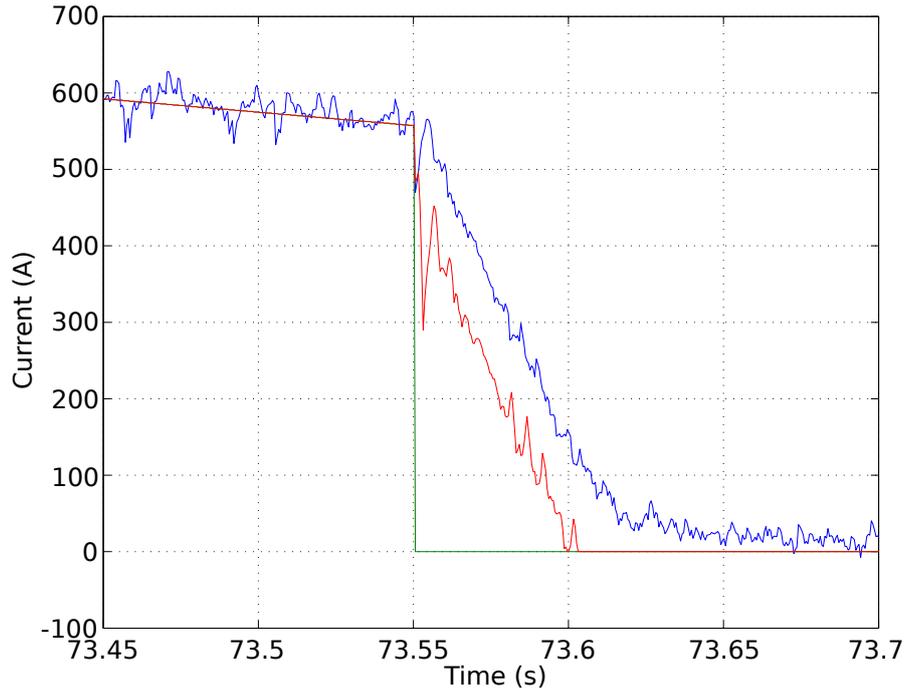


Figure 5.12: Slew rate - (blue) estimated current; (green) adjusted reference and (red) slope limited reference

Proportional Integral Derivative control

The PID [Fran 91] is, together with the bang-bang¹ type of control [Sonn 64], certainly the most widely used control scheme. This mechanism is based on the difference between the desired value and the present value of the controlled parameter. This difference is commonly referred to as the *error* signal. In the case of the present application the error is given by $e_k = I_{k+1}^{ref} - I_k^+$, i.e., the difference between the desired value of the current for the next control cycle and its present value. In its general form, the control action V_k^{req} , i.e. the request to the VCVS, is calculated using (5.13) where K_p , K_i and K_d are design parameters. The control action is therefore based on a linear combination of the error, its integral and its derivative. It is expected that the effect of the control action steers the system in the direction of reducing $|e_k|$ however, that may not always be the case.

$$V_k^{req} = K_p \times e_k + \frac{K_i}{F_s} \times \sum_{i=1}^k \frac{e_i + e_{i-1}}{2} + K_d F_s \times (e_k - e_{k-1}) \quad (5.13)$$

Although the PID is a well established and thoroughly studied method, typically credited for its robustness and for its obliviousness of the nature of the underlying process, it is hardly exempt of weaknesses [Athe 99]. In particular, the fact that it doesn't take into account information about the underlying process means that, rather than predicting the system trajectory in state-space and anticipating the control action, the PID has to *wait* for the system to deviate from the desired state in order to start taking some action. It is therefore intuitive to realize that, especially under non-stationary conditions, any system output controlled with a PID is always lagging behind its target state. Also, because the control action's calculation is solely based on a linear combination (with

¹Bang-bang control, also known as on-off control, is a scheme for which typically the control action is either disabled or at a fixed *intensity* level.

fixed coefficients²) of linear operators acting on the error signal, means that it cannot be the solution of an optimization problem in the entire operational domain. That is to say that PID control is not optimal.

In the present context, the use of external magnetic perturbations to measure the RFA effect on JET plasmas requires operation of the EFCCs at various frequencies within the system's capabilities, typically in the range 0–50 Hz. PID-based control is incapable of providing optimal control across this whole range of frequencies. In the present application one has chosen to, while giving the user the option of using a basic PID algorithm, develop a more sophisticated model-based approach in an attempt to overcome some of the aforementioned PID weaknesses.

Model Predictive Control

The model-based method chosen was MPC. The main principle of MPC [Mora 99], also known as receding horizon control, is to solve the open-loop control problem, in real-time and on every control cycle for a finite temporal horizon using the present plant state as an initial condition, in order to determine the sequence of control actions optimizing a given criterion. Although the optimal set of control actions, at a given time instant, is calculated for a finite prediction horizon, only the control action for the present is actually executed. Whenever possible, the most relevant physical limitations of the system, such as actuator bounds and state domain, are taken into account so that the solutions of the optimization process are constrained to a realistic operational space.

Clearly an MPC implementation can be rather demanding in terms of computational burden. Keeping in mind that each optimization process needs to be performed and complete within two consecutive control cycles, the choice of the prediction horizon is tightly coupled with the system's bandwidth specifications/requirements which, in turn, strongly influence the choice of the control period. One of the greatest strengths of MPC-based control is that it solves the optimal control problem for a finite prediction horizon, in real-time, using the present plant state as an initial condition rather than, as is done in conventional control approaches, pre-computing off-line a control law (even if optimal with respect to some criterion) over all possible system states.

The optimal control problem is expressed as a cost function minimization problem. A typical form for the MPC cost function is given by (5.14), where N_p is the number of discrete time instants in the prediction horizon and $\|\cdot\|$ denotes an unspecified norm. \mathbf{u}_k is the set of control variables i.e. the physical quantities used to control the system and, obviously $\Delta\mathbf{u}_k$ is their change between two consecutive control cycles. $\mathbf{p}_k^{pred} = \mathbf{p}_k^{pred}(\mathbf{p}_{k-1}^{pred}, \mathbf{u}_k)$ is the set of process variables being controlled (obviously a function of their starting value \mathbf{p}_{k-1}^{pred} and the inputs \mathbf{u}_k) while \mathbf{r}_k denotes their target values. It is worth noting that, for $i = 1$, $\mathbf{p}_k^{pred} = \mathbf{p}_k^{pred}(\mathbf{p}_{k-1}^{est}, \mathbf{u}_k)$, i.e., the initial condition is the estimate rather than the prediction of the process variable set in the previous control cycle. Finally, w_p , w_u and $w_{\Delta u}$ are simple weights. The leftmost term on the right-hand side of (5.14) penalizes deviations of the variables under control from their target values. The middle term penalizes the *intensity* of the control action while the rightmost term penalizes the excursions in consecutive control actions. While minimizing the leftmost term is actually the ultimate goal of the entire optimization process, the other terms provide soft constraints that often help mitigate practical problems related to the control effort. It is worth noting that, in case $\|\cdot\|$ is implemented as the *2-norm* and $\mathbf{p}_k^{pred}(\mathbf{p}_{k-1}^{pred}, \mathbf{u}_{k-1})$ is linear, minimizing J_k is

²Gain scheduling is certainly possible but assumes some knowledge about the process.

a linear least-square minimization problem which is analytically solvable.

$$J_k = \sum_{i=1}^{N_p} \left[w_p \|\mathbf{p}_{k+i}^{pred}(\mathbf{p}_{k+i-1}^{pred}, \mathbf{u}_{k+i-1}) - \mathbf{r}_{k+i}\| + w_u \|\mathbf{u}_{k+i-1}\| + w_{\Delta u} \|\Delta \mathbf{u}_{k+i-1}\| \right] \quad (5.14)$$

Let one assume that the cost function J_k one wishes to minimize in every control cycle k is given by (5.15). The aim is therefore to calculate the voltage requests, for the present (V_k^{req}) and the subsequent control cycle (V_{k+1}^{req}), that minimizes J_k . Only V_k^{req} , the voltage request performed by the controller at the end of the present control cycle k , is actually performed. In this case, the prediction horizon is $N_p = 2$. Having specified the desired bandwidth of the system to be 100 Hz, it is reasonable to have a control application running at a rate which is 10 – 20 times faster. Therefore, the requirement of a 2 kHz control cycle means that performing the online optimization within the application's cycle time with $N_p > 2$ was not possible with the available hardware.

$$J_k = \left(I_{k+1}^{pred} - I_{k+1}^{ref} \right)^2 + \left(I_{k+2}^{pred} - I_{k+2}^{ref} \right)^2 + \mu \left(V_k^{req} - V_{k-1}^{req} \right)^2 + \varepsilon \left(V_{k+1}^{req} - V_k^{req} \right)^2 \quad (5.15)$$

In this case there is only one process variable under control which is the estimate of the current (I_k^+) in the EFCC circuit as given by the KF, see (5.4). In the light of the previous discussion one has that $I_{k+1}^{pred} = I_{k+1}^{pred}(I_k^+, V_k^{req})$ and $I_{k+2}^{pred} = I_{k+2}^{pred}(I_{k+1}^{pred}, V_{k+1}^{req})$. These will therefore need to be obtained and substituted inside (5.15). It is worth clarifying that I_{k+1}^{pred} is a prediction, performed during control cycle k , of what the current in the circuit will be at the beginning of control cycle $k + 1$, that is, after the voltage request V_k^{req} at the end of control cycle k has been performed.

In order to proceed, let one now consider the behaviour of the VCVS. The VCVS, being a non-perfect voltage amplifier, will have a finite complex transfer function associated with it. This is to say that the VCVS does not instantaneously provide the requested voltage. Let one specify a linear model of the form given by (5.16) to express the VCVS's output voltage V_k , at time index k , as a linear combination of past output voltage measurements and the VCVS's readback¹ of past voltage requests V_k^{rbk} . Two assumptions are implicit here: (1) intuition-based assumption that the VCVS's output voltage depends not only on the present and past voltage requests but also on the past output voltages and (2) these are linear dependencies. Note that the two most recent output voltage measurements are assumed to be unknown. The one referring to $i = 0$ (b_0) is obvious since the present output voltage is on the left-hand side of (5.16). The absence of the b_1 coefficient is because, as will become clear shortly below, if one uses (5.16) to calculate V_{k+2}^{pred} , the output voltage at time index $k + 1$ (V_{k+1}) will not be known/available yet².

$$V_k = \sum_{i=2}^{N_v} b_i V_{k-i} + \sum_{i=0}^{N_r} c_i V_{k-i}^{rbk} \quad (5.16)$$

If one now assumes that $V_k^{req} \approx V_{k+1}^{rbk}$ (i.e. the voltage request calculated and performed at the end of control cycle k is approximately the voltage request acknowledged by the

¹ V_k^{rbk} comes originally from an analog signal, produced by the VCVS, containing the last voltage request acknowledged by the VCVS itself. This signal is subsequently digitized in the controller application. In noise-less environments, the voltage request performed at the end of control cycle k (V_k^{req}) is exactly the same as the readback voltage at the beginning of control cycle $k + 1$ (V_{k+1}^{rbk}).

²Note that this predictive calculation is performed in the present, i.e. in control cycle k .

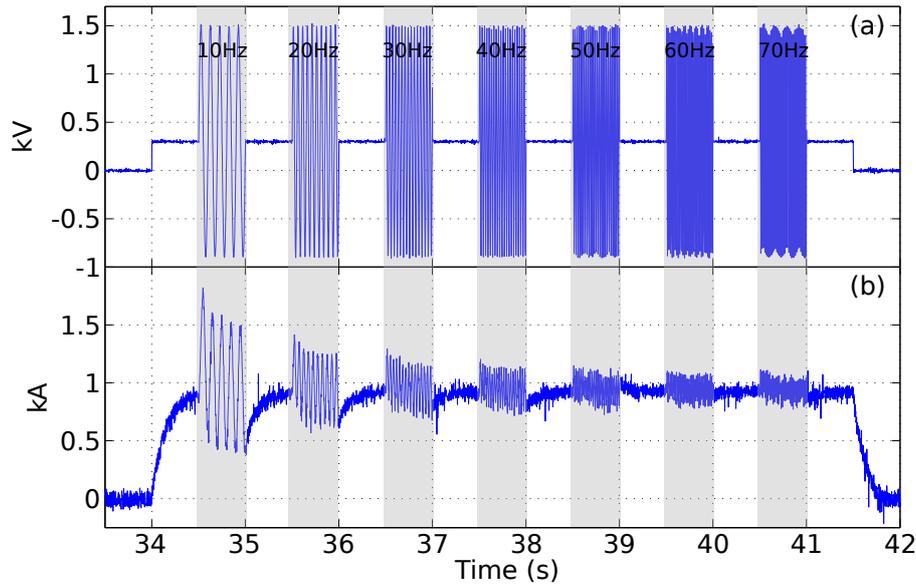


Figure 5.13: Pre-configured voltage control - (a) Readback of the voltage request acknowledged by the VCVS and (b) Current in the EFCC circuit.

VCVS at the beginning of control cycle $k + 1$), and therefore that $V_{k+1}^{req} \approx V_{k+2}^{rbk}$ then, using (5.16), yields (5.17). V_{k+1}^{pred} and V_{k+2}^{pred} are the VCVS's output voltage predictions for the future (for time indexes $k + 1$ and $k + 2$, respectively). It is worth stressing out that these predictions are performed in the present, i.e., at time index k .

$$\begin{cases} V_{k+1}^{pred} = \sum_{i=1}^{N_v-1} b_{i+1}V_{k-i} + c_0V_k^{req} + \sum_{i=1}^{N_r} c_iV_{k-i}^{rbk} \\ V_{k+2}^{pred} = \sum_{i=0}^{N_v-2} b_{i+2}V_{k-i} + c_0V_{k+1}^{req} + c_1V_k^{req} + \sum_{i=2}^{N_r} c_iV_{k-i}^{rbk} \end{cases} \quad (5.17)$$

System identification techniques are required to determine the optimal set of coefficients b_i and c_i for reproducing the VCVS's behaviour. In order to do this, a frequency scan (0 – 70 Hz in 10 Hz steps) was performed where the VCVS was driven in pre-programmed voltage control, see figure 5.13, and the coefficients were subsequently determined from the acquired data using Matlab's auto-regressive exogenous modelling function. The reported accuracy of fit was 86% and the coefficients obtained are depicted in figure 5.14. It was observed that, for N_v and $N_r > 11$, the improvement in the accuracy of the fit was marginal. It is evident that, in overall terms, the voltage request has a bigger impact, especially in the first four time lags, than the set of past output voltages.

Basic VCVS tests revealed that it was in fact unable to provide an absolute output voltage larger than 1.8 kV. Figure 5.15 shows the operational domain of the amplifier in the context of an $N_p = 2$ prediction horizon. In the same figure, the red labels numbered 1 – 8 designate a particular zone in the frontier of the domain and 0 its enclosure. These will be used later on as a reference for the constrained minimization of the cost function.

In order to finalize the derivation of the MPC-based control equations for the problem at hand, let one start from the general RL circuit equation (5.2) and rewrite it in the form given by (5.18), very similar to the KF's prior current estimate (5.3). I_k is the current flowing in the EFCC circuit, V_k is the VCVS's output voltage, $\alpha = L_{eq}F_s$ and $\beta =$

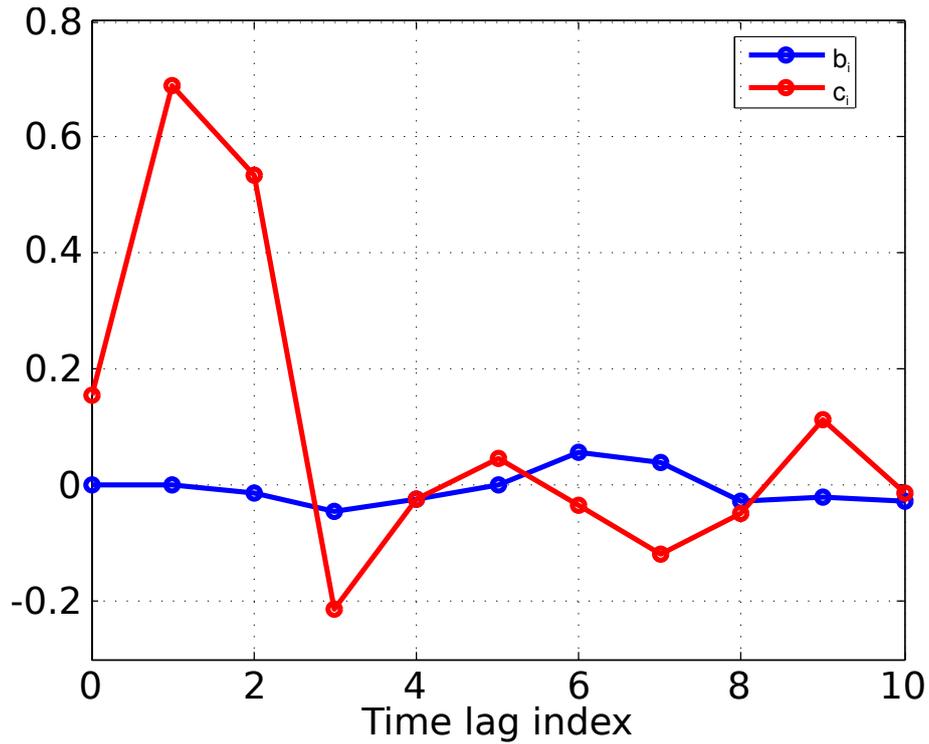


Figure 5.14: VCVS linear model coefficients.

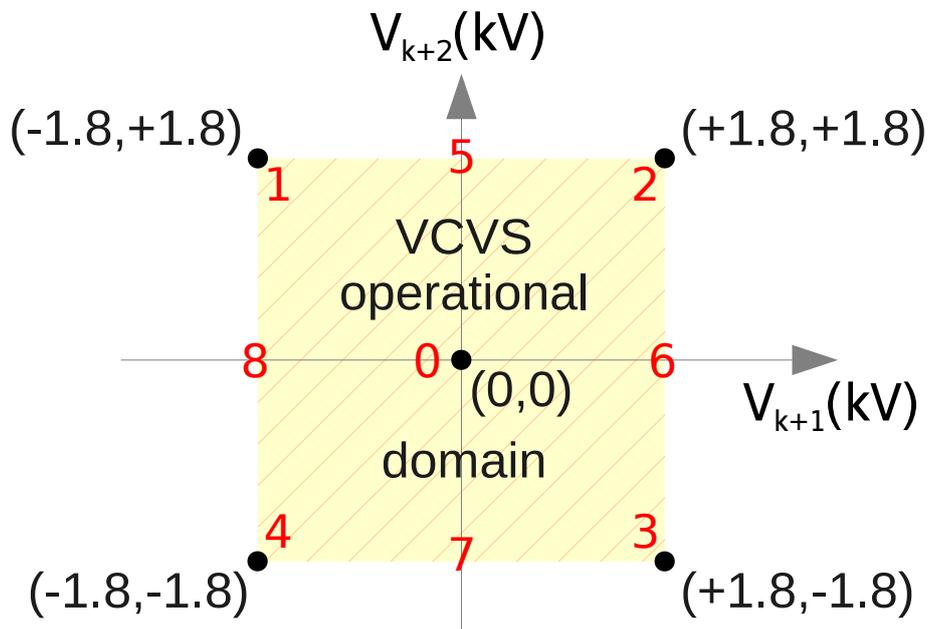


Figure 5.15: VCVS linear operational domain - red indexes are labels to identify the regions in the frontier of the operational domain (1 – 8) and inside the operational domain (0).

$1/(R_{eq} + \alpha)$. As previously discussed R_{eq} , L_{eq} and F_s are, respectively, the equivalent resistance, equivalent inductance and sampling frequency. The mutual inductances are assumed to be zero.

$$I_k = \alpha\beta I_{k-1} + \beta V_k \quad (5.18)$$

In the light of the previous, and in order to obtain the current predictions I_{k+1}^{pred} and I_{k+2}^{pred} for the next two control cycles one writes (5.19). It is worthwhile noting that for the first prediction the starting point is the KF's present current estimate I_k^+ while for the second prediction the starting point is the previous one.

$$\begin{cases} I_{k+1}^{pred} = \alpha\beta I_k^+ + \beta V_{k+1}^{pred} \\ I_{k+2}^{pred} = \alpha\beta I_{k+1}^{pred} + \beta V_{k+2}^{pred} \end{cases} \quad (5.19)$$

So, substituting (5.17) in (5.19) and subsequently the latter in (5.15) yields $J_k = J_k(V_k^{req}, V_{k+1}^{req})$, where all other quantities are either available measurements/estimates such as I_k^+ , V_{k-i} and V_{k-i}^{rbk} , or known parameters such as R_{eq} , L_{eq} and F_s . It is straightforward to verify that the minimization of this expanded version of (5.15) expresses a convex problem hence admitting only one minimum which is the necessarily the global.

In order to find V_k^{req} and V_{k+1}^{req} that provides the unconstrained minimum of J_k one has to solve, as usual, $\nabla J(V_k^{req}, V_{k+1}^{req}) = \vec{0}$ yielding (5.20), where $\psi = \alpha\beta c_0 + c_1$.

$$\begin{pmatrix} \beta^2(c_0^2 + \psi^2) + \mu + \xi & \beta^2 c_0 \psi - \xi \\ \beta^2 c_0 \psi - \xi & \beta^2 c_0^2 + \xi \end{pmatrix} \begin{pmatrix} V_k^{req} \\ V_{k+1}^{req} \end{pmatrix} = \begin{pmatrix} f(I_k^+, I_{k+1}^{ref}, I_{k+2}^{ref}, V_k^{req}, V_{k-i}, V_{k-i}^{rbk}) \\ g(I_k^+, I_{k+2}^{ref}, V_k^{req}, V_{k-i}, V_{k-i}^{rbk}) \end{pmatrix} \quad (5.20)$$

The previous is of the form $\mathbf{\Lambda} \times \mathbf{V}^{req} = \mathbf{\Gamma}$. The square matrix $\mathbf{\Lambda}$ admits an inverse if (5.21) is satisfied. It is straightforward to note that: if R_{eq} , L_{eq} , μ , ξ , c_0 and c_1 are all positive, then (5.21) is always satisfied.

$$\det(\mathbf{\Lambda}) = \beta^2 (\beta^2 c_0^4 + 2c_0^2 \xi + \psi^2 \xi + \mu c_0^2 + 2c_0 \psi) + \mu \neq 0 \quad (5.21)$$

Under these assumptions, the solution of the unconstrained minimization problem is always given by $\mathbf{V}^{req} = \mathbf{\Lambda}^{-1} \times \mathbf{\Gamma}$. Although by itself this is, in principle, already an improvement over PID-based control, there is still the VCVS's operational domain to take into account in the minimization process. Strictly speaking it becomes a linear least-squares minimization problem with inequality constraints. Although in principle solvable using Lagrange multipliers or more sophisticated quadratic programming methods, the approach taken herein was far simpler. First the unconstrained minimization problem is solved and, should the solution \mathbf{V}^{req} lie inside the VCVS's operational domain, no further processing is required. Should it lie beyond such domain, then the frontier is analytically scanned to find out the values of V_k^{req} and V_{k+1}^{req} for which J_k is minimal.

In the following section, preliminary results are shown regarding the VCVS modelling and controller performance.

5.4 Preliminary results

In this section, preliminary results of the operation of the system are presented. In particular, an example is shown of mixed operation of the controller system interleaving pre-configured current references with the ones provided by RTCC in real-time thus commissioning the reference multiplexing logic.

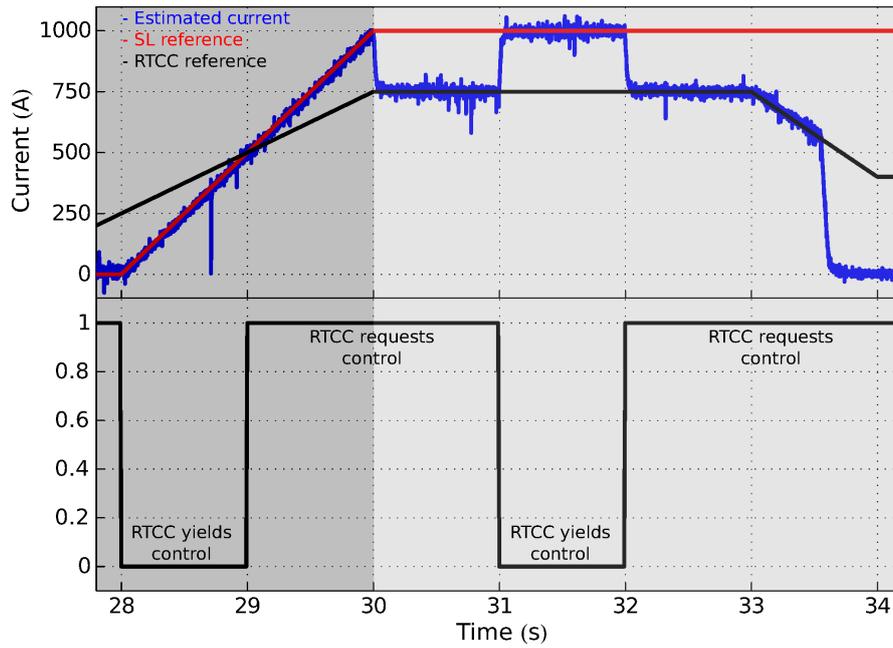


Figure 5.16: Commissioning of the reference multiplexing logic.

Also, a basic assessment of the quality of the VCVS modelling and controller performance is done in comparative terms for both the PID and MPC-based control modes. Regarding MPC control, also different configurations of μ and ξ parameters were tested.

5.4.1 Current reference multiplexing logic assessment

This section illustrates the commissioning tests of the reference multiplexing logic, see figure 5.16. The blue curve represents the KF posterior estimate I_k^+ of the current flowing in the EFCC circuit, the red curve is the pre-configured SL reference and the black curve is the real-time reference computed at a rate of 500 Hz by RTCC and sent to the controller system over JET's ATM-based RTDN. The dark and light gray areas are pre-configured by the SL, respectively, as the time windows for which control from RTCC isn't and is allowed. Furthermore, even if in control, RTCC can yield it by instructing the controller to revert the current references used for control to the ones pre-configured by the SL. So, analysing figure 5.16 one notes that, for $t < 28$ s, RTCC is requesting control but the controller is following the SL reference. Then, for $28 < t < 29$ s, RTCC refrains its control request. For $29 < t < 31$ s RTCC requests control and is finally able to get it (notice that the current starts to follow the black reference rather than the red one) starting from $t = 30$ s corresponding to the beginning of the allowed window. Further ahead, for $31 < t < 32$ s and although inside the window that permits RTCC control, RTCC yields control and the current goes back to following the SL(red) reference. Finally, for $32 < t < 33.5$ s, RTCC claims back control and the current starts tracking the real-time reference once more. At $t = 33.5$ s, a stop test is performed to evaluate the responsiveness of the system under those circumstances.

This test successfully completes the commissioning of the reference multiplexing feature.

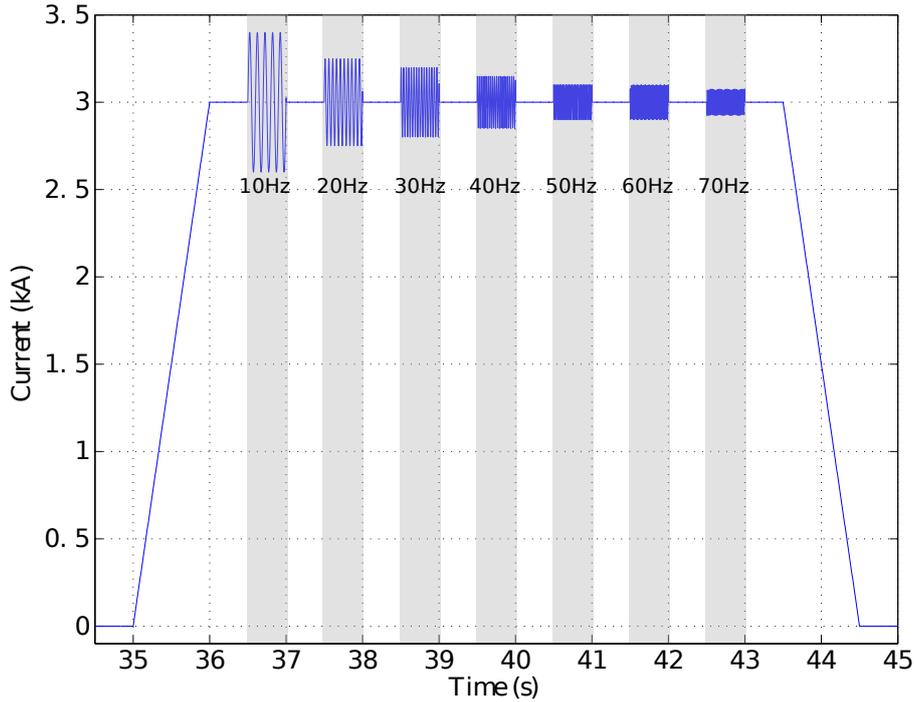


Figure 5.17: Pre-configured current reference.

5.4.2 VCVS model performance

In this section, the quality of the VCVS modelling is succinctly analyzed. The results presented herein were obtained at an early stage of the development of the system, for which the control cycle had not yet been optimized to run at the 2 kHz rate. The control cycle of the controller application in the following tests was 500 Hz.

A current reference waveform consisting of a 3 kA baseline with seven 0.5 s time windows of AC operation with a frequency range of 10 – 70 Hz in steps of 10 Hz was pre-configured, see figure 5.17. The amplitude of each AC portion was chosen not to exceed the dI/dt limits so that there was no chance of tripping the toroidal field circuit’s protection system, see section 5.3.3. Since $dI/dt \propto \omega \times A$, where ω is the AC frequency and A is its amplitude, an increase in ω needs to be compensated by a decrease in A in order to limit dI/dt .

The idea here is basically to assess the accuracy of the output voltage prediction V_{k+1}^{pred} which is based on the VCVS model. The VCVS model, as previously discussed, is a linear model given by (5.16) and for which the coefficients b_i and c_i were calculated from the data obtained in an experiment performed in pre-programmed voltage control (hence without feedback) as shown in figure 5.13.

Four different controller configurations were used to control the current in the EFCC circuit using the reference depicted in figure 5.18. The first three controller configurations are MPC-based and characterised by $\mu = \xi = 10^{-4}$, $\mu = \xi = 10^{-3}$ and $\mu = \xi = 10^{-2}$, respectively. The fourth controller configuration is PID-based with $K_p = 5$, $K_i = 20$, $K_d = 0$. A set of relevant time traces obtained (within the 30 Hz time window) are shown in figure 5.18 for the different controller configurations. The first noticeable thing is the saturation of the voltage request in all configurations. The reason for this is that the DC (resistive) component alone consumes almost 1 kV of the total VCVS’s available output voltage (1.8 kV) thus essentially leaving $\sim 800V$ for the AC’s positive half-cycles (mainly

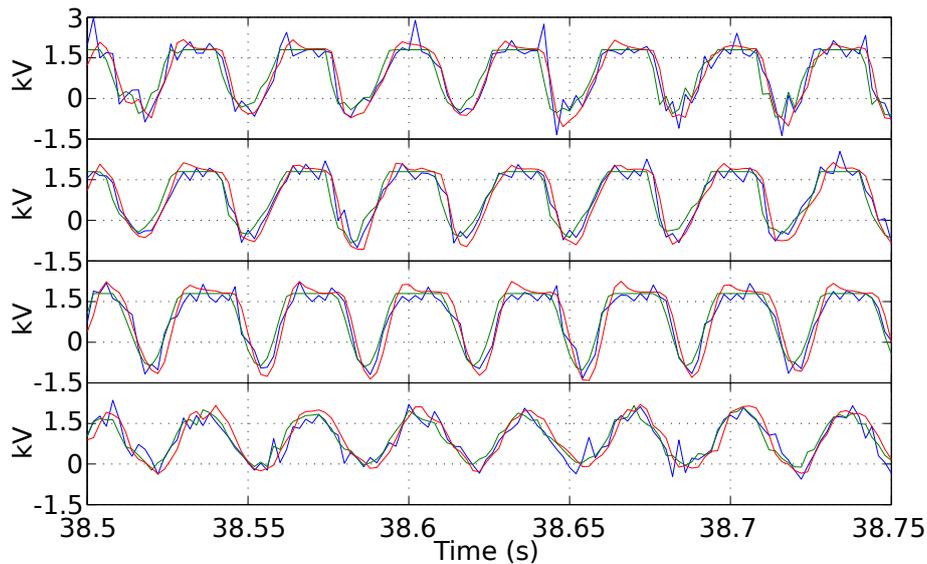


Figure 5.18: VCVS model assessment - (blue) VCVS output; (green) Controller request; (red) Predicted VCVS output - (a) MPC with $\mu = \xi = 10^{-4}$, (b) MPC with $\mu = \xi = 10^{-3}$, (c) MPC with $\mu = \xi = 10^{-2}$, (d) PID with $K_p = 5, K_i = 20, K_d = 0$.

inductive). The second observation is that, in general, there is a lag between the controller request and the VCVS's output voltage thus confirming the need for modelling the VCVS's transfer function. A third remark is that, in general, there is a visibly acceptable agreement between the model's predictions and the true obtained voltage for which a quantification is presented in figure 5.19. By performing the DFT of the output voltage and the predicted output voltage, one is able to plot the difference in their amplitude normalized to the amplitude of the output voltage, i.e. $[|DFT(V_k^{pred})| - |DFT(V_k)|]/|DFT(V_k)|$, as shown in figure 5.19(a) for each of the *scanned* frequencies. The immediate conclusion is that the model tends to overestimate the amplitude of an oscillatory output voltage. The exact cause of this effect is not yet at this time. By performing the cross-spectrum analysis of the output voltage and its prediction, one is able to estimate their phase delay normalized to the oscillation period and plot it as shown in figure 5.19(b). From this it can be seen that the delay in the estimation of the VCVS's output voltage is never larger than 10% of the period of oscillation.

5.4.3 Controller performance

In this section, the controller performance is succinctly assessed. Just like in the previous section, the results presented herein were obtained at an early stage of the development of the system, for which the control cycle had not yet been optimized to run at the 2 kHz rate. The control cycle of the controller application in the following tests was 500 Hz.

The same different controller configurations used in the previous section were also used herein and an example of relevant time traces obtained (within the 30 Hz time window) are shown in figure 5.20. It is not surprising to observe that, in figure 5.20(a), for the AC's negative half-cycle both the blue and green curves adequately match the reference whereas for the positive half-cycles that is not the case for any of the curves. The fact that the MPC configurations with $\mu = \xi < 10^{-3}$ are able to provide better control (closer to the target reference) than with the other configurations can be seen from the voltage

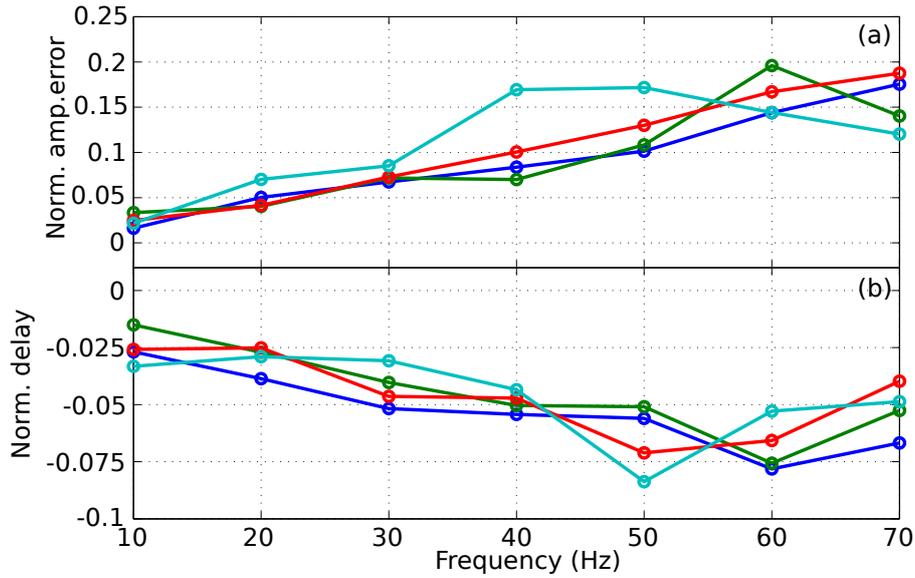


Figure 5.19: VCVS model assessment - (blue) MPC with $\mu = \xi = 10^{-4}$, (green) MPC with $\mu = \xi = 10^{-3}$, (red) MPC with $\mu = \xi = 10^{-2}$, (cyan) PID with $K_p = 5, K_i = 20, K_d = 0$ - (a) Normalized amplitude error; (b) Delay normalized to period

requests shown in figure 5.20(b). It is clear that these MPC configurations react faster than the others. This suggests that, for instance, the penalizing factors $\mu = \xi = 10^{-2}$ of the other MPC controller configuration severely limits the output bandwidth of the control action by restricting voltage request excursions in consecutive cycles. Regarding the PID-based control case, its output bandwidth would improve by increasing K_p and eventually decreasing K_i however, ultimately, it is never optimal. The previous discussion is also confirmed by figure 5.20(c) where the value of the cost function J_k is shown. The fact that, for $\mu = \xi = 10^{-2}$, its baseline is at least one order of magnitude above the others and that its normalized excursions are smaller than the others, indicates that the cost function minimization process is dominated by the terms penalizing excursions in consecutive voltage requests. Finally figure 5.20(d) shows, with respect to the VCVS's operational domain (refer to figure 5.15), where the solutions of the constrained minimization of J_k lie. As expected, in the positive AC half-cycles the solutions lie mainly in the frontier of the operational domain whereas in the negative half-cycles they lie inside the operational domain (meaning that the constrained and unconstrained minimization solutions coincide).

In order to further quantify the performance of the different control schemes under scrutiny, the exact same study that was done to assess the VCVS model in the previous section was performed to assess the control performance. By performing the DFT of the target reference current and the estimated current obtained, one is able to plot the difference in their amplitude normalized to the amplitude of the reference current, i.e. $||DFT(I_k^{est})| - |DFT(I_k^{ref})|| / |DFT(I_k^{ref})|$, as shown in figure 5.21(a) for each of the scanned frequencies. Whereas all MPC-based control schemes exhibit similar performance, the PID substantially attenuates the amplitude of the oscillatory output current with respect to the reference. By performing the cross-spectrum analysis of the current reference and the estimated current, one is able to infer their phase delay normalized to the oscillation period and plot it as shown in figure 5.21(b). It can be seen that, with respect to the delay, on average the MPC-based schemes with $\mu = \xi \leq 10^{-3}$ (blue and green

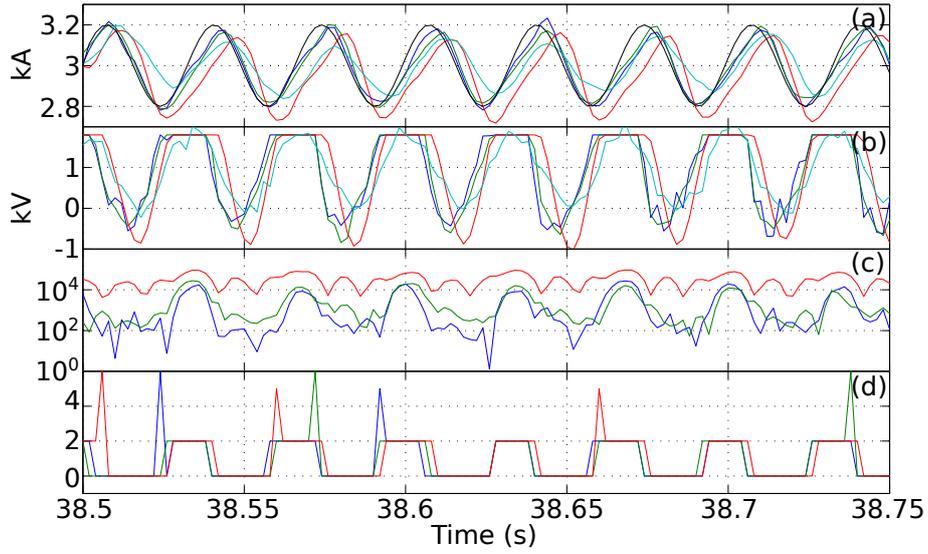


Figure 5.20: Control performance - (blue) MPC with $\mu = \xi = 10^{-4}$, (green) MPC with $\mu = \xi = 10^{-3}$, (red) MPC with $\mu = \xi = 10^{-2}$, (cyan) PID with $K_p = 5, K_i = 20, K_d = 0$ and (black) reference - (a) Estimated current; (b) Requested voltage; (c) Cost function; (d) VCVS operational space index (refer to figure 5.15).

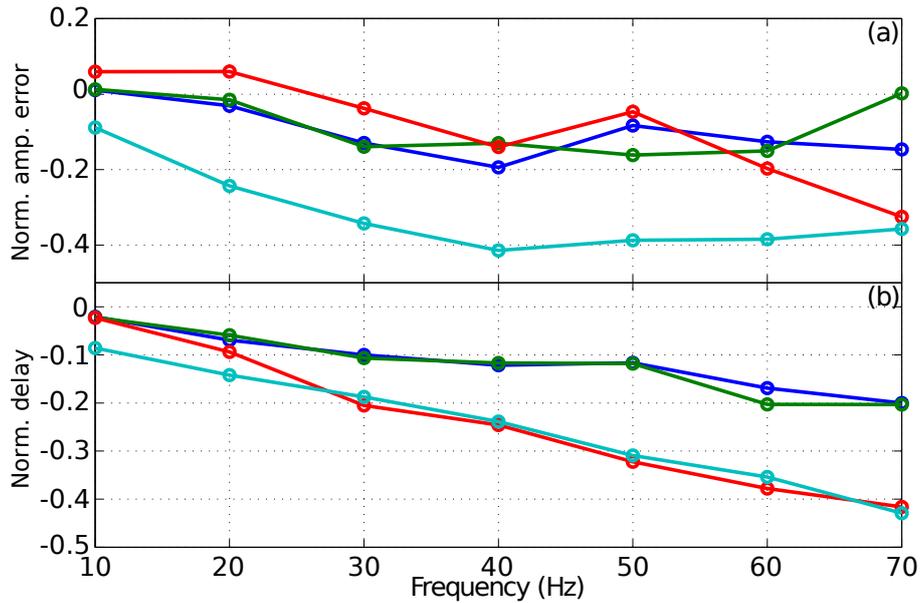


Figure 5.21: Control performance - (blue) MPC with $\mu = \xi = 10^{-4}$, (green) MPC with $\mu = \xi = 10^{-3}$, (red) MPC with $\mu = \xi = 10^{-2}$, (cyan) PID with $K_p = 5, K_i = 20, K_d = 0$ - (a) Normalized amplitude error; (b) Delay normalized to period

curves) behave very similarly. Likewise, both the MPC-based approach with $\mu = \xi = 10^{-2}$ and the PID-based control behave very similarly.

Chapter 6

Final remarks and future work

This dissertation presented a number of KF-based applications suitable for real-time signal processing in tokamak fusion research and evidenced how robust and adaptable Kalman filter methods are thus demonstrating their potential to this day, more than 50 years after being first proposed.

The first application presented herein employs a linear KF, used in the context of the real-time signal processing of the MSE diagnostic at JET as an alternative to the standard Fourier-based lock-in implementations. It is shown to be an efficient and flexible lock-in implementation for the real-time estimation of the amplitude of quasi-stationary harmonic components. At its core is the KFHE, designed to estimate, simultaneously, both in-phase and quadrature components of a sum of quasi-periodic signals immersed in white noise. Its implementation is relatively straightforward and relies on a state transition model that evolves, simultaneously, the phase of each component assuming a constant amplitude. Subsequently, measurement information is incorporated to update the model-based estimates. The proposed approach is, in general, computationally heavier (depending on the number of frequency components being simultaneously estimated) than Fourier-based lock-in amplifier implementations but its advantages clearly make up for the added computational burden in some circumstances. While for offline analysis the Fourier-based methods are valuable, in real-time implementations their performance may be limited. For example, for output rates of the order of the raw data's input rate, Fourier methods are not able to provide adequate filtering since the decimation process is tied to the frequency response of a square window of a limited amount of samples. On the contrary, KF-based methods, even if using the same square window for decimation purposes, always provide some degree of filtering due to their state transition model thus attenuating all quasi-periodic components not included therein. KF-based methods are clearly real-time native techniques in the sense that they're data stream oriented rather than a data block oriented. The proposed KF-based lock-in amplifier implementation was used for the real-time amplitude estimation of the quasi-periodic APD components of experimental MSE signals. This technique has demonstrated to have some advantages when compared to the ordinary Fourier-based methods in issues such as SNR, sensitivity to phase/frequency jitter or mitigation of the biasing effect of ELMs in the MSE measurements. The proposed method has been found to be a valuable contribution of the overall quality improvement of MSE data for usage in plasma equilibrium reconstruction schemes.

A KF-based method was also developed for the real-time analysis of coherent magnetic structures in tokamak research. The method was applied to magnetic sensor data

from typical ISTTOK plasma pulses revealing, for pulse #17081, an $m = 2$ observed at ~ 160 kHz. In particular, the method was shown to be especially advantageous in cases where the events are so fast that only a very limited amount of samples is covering the event duration. Although temporal averaging is still required as a replacement for the (statistically rigorous but unpractical) ensemble average required by the coherence analysis, the streaming nature of the filter means that, unlike the Fourier-based approach, the frequency resolution is not tied to the number of signal samples used. In a comparative study, it was evident that the Fourier-based method struggles to decorrelate spatially incoherent fluctuations due to the lack of samples to increase its effective frequency resolution.

Another KF-based method was developed for the synchronous detection of resonant plasma events associated with the TAE antenna diagnostic at JET. Although in this case (as opposed to the signal processing of MSE data) there is only one tone of interest, its frequency is now non-stationary and its evolution is known in real-time. Two KF-based approaches have been developed. The first approach utilizes a non-stationary state transition model in a linear KFHE implementation where the phase advance is tied to the frequency evolution of the excitation signal. The second approach utilizes the FT configuration in an EKF implementation which not only dispenses the knowledge of the exact frequency evolution, but also estimates it. The fact that the EKFFT dispenses the knowledge of the excitation waveform's frequency evolution is shown to have little impact on its amplitude estimation performance when compared to the linear KF. These KF-based approaches are also very flexible and low maintenance requiring only the tuning of the ratio between the covariance of the noise associated with the measurements and the uncertainty in the underlying state transition model.

Three different non-linear Kalman filter-based algorithms were successfully applied to the real-time frequency tracking and harmonic estimation of dominant MHD activity in JET pulses. These methods are inherently real-time native and able to provide accurate estimates at fast sampling rates without requiring any sort of temporal averaging mechanisms. In particular, the single-tone adaptive configuration developed in section 4.4.2 for the estimation of the dominant component in JET magnetic signals has proven to be extremely reliable and robust even in the presence of unmodelled disturbances such as ELMS. In fact, the main requirement for the reliable operation of these algorithms is the presence of a signal component with an amplitude $\gtrsim 2$ times larger than the second strongest one (and the noise level) at any given time instant so that it can be unambiguously identified. This covers the overwhelming majority of cases regarding MHD activity in JET pulses. When performing multi-tone estimation though, extra computational effort and extra requirements are necessary especially regarding the minimum inter-frequency separation allowed for the estimates thus adding to the complexity.

The filter configurations used in both single-tone and multi-tone adaptive implementations herein have demonstrated to provide accurate estimates for a set of > 20 plasma pulses exhibiting notorious MHD activity. Such set encompasses a number of plasma configurations with the following parameter ranges for current (1.4 – 3.0 MA), density ($6.4 - 22.9 \times 10^{19} \text{ m}^{-2}$), toroidal magnetic field (2.2 – 2.8 T) and external input power (11.2 – 26 MW).

In the case when only the dominant signal component is being estimated, the adaptive mechanism has proven irrefutably superior to the non-adaptive one not only in terms of tracking capability but also in terms of rejecting inter-frequency contamination. Under these conditions, and for a similar level of complexity and computational demand, the

EKFFT and rEKFFT present similar performance. Mathematically, the rEKFFT requires two 2×2 matrix inversions for the calculation of the prior covariance matrix, equation (2.80), while the EKFFT requires none. Extra care is therefore required for the real-time implementation of the rEKFFT in order to ensure the stability of the algorithm itself. On the other hand, and although not particularly evident in the examples presented here, the rEKFFT is more resilient to disturbances than the EKFFT. The UKFFT consistently provides better estimates and better inter-frequency rejection performance on average than both the EKFFT and rEKFFT however, the algorithm itself requires one 2×2 matrix inversion and two 2×2 Cholesky decompositions hence also requiring some care in its real-time implementation.

In the case of the *three-independent-tone* estimation, both the rEKFFT and UKFFT have demonstrated to be superior to the EKFFT. Matrix inversions and Cholesky decomposition requirements are the same only now matrix dimensions are 9×9 . There should a strong motivation to perform multi-tone rather than single-tone estimation as the added complexity, computational burden and non-guaranteed accuracy in non-dominant component tracking are certainly handicaps to take into account. However, as observed in the examples disclosed herein, good estimation performance of non-dominant tones can be achieved and the individual residues always provide a reliable validation measure of the tracking efficacy. Furthermore, simultaneous multi-tone estimation in principle allows for cleaner estimates as the components being tracked will not propagate, via the innovation, into each others' posterior state estimate.

In addition to the previous, a novel method based on linear and non-linear KFs aimed at the real-time mode number estimation of magnetic perturbations in magnetically confined fusion plasmas was proposed and results of its application to experimental data of the JET tokamak was presented. The proposed method implements an inherently real-time approach as it acts on sample streams rather than sample blocks unlike, for instance, the coherence spectrum or SVD based methods. The time resolution of this method is essentially the raw data's sampling rate. In the case of the coherence spectrum or SVD based methods, although their temporal resolution is controlled by the amount of raw data overlap used in the analysis, their time localisation will be disperse over the block of samples being processed unlike the KF based approach.

Contrary to Fourier based methods, frequency is treated as a continuous quantity by the EKFFT algorithm. Fourier based methods employ a fixed array of frequency bins whose separation is determined by the sampling frequency and the number of samples used in the analysis whenever the popular FFT algorithm is used. SVD, by itself, does not provide frequency information directly thus relying on additional method(s), such as Fourier, to extract the frequency information from the temporal decomposition. The end-to-end KF based method utilises a fixed array of mode numbers for simultaneous estimation and is able to provide the frequency of the dominant signal component as well as the amplitudes of each mode. It is worthwhile mentioning that, although not shown herein, one of the by-products of the analysis presented in this work is the determination of the time dependent phase evolution of each mode number given the KFMP's estimates of in-phase and quadratures of all modes. This time dependent phase of the MHD activity, together with a spatial characterization of the mode, may be required in feedback stabilization mechanisms. In particular, for the case of feedback stabilization using modulated ECRH and ECCD on NTMs [Zohm 07], it provides the pivotal indication of when the Radio Frequency (RF) antenna is facing an island's *O-point* or *X-point* and thus when the rf waves should be launched. The conceptual design of the KF mode

identification presented in this work is solely focused on the tracking of the time dependent frequency/phase of the mode and the identification of the toroidal eigenvalues of the associated magnetic oscillations. An extension of the design to address also the control of the modes by RF power would require the addition of dedicated modules for describing the non-linear interaction of the mode with the RF waves.

The novel method proposed herein has been demonstrated to be able to resolve mode number degeneracy observed at a given non-stationary frequency, see sections 4.4.3 and 4.4.3, and to provide accurate and robust amplitude estimations in the presence of unmodelled disturbances such as ELMs and sawtooth crash events. In fact, and unlike all other Fourier and SVD based methods, the KF's weighting of state prediction (for which the starting point is the previous posterior state estimation) and measurements, based on their uncertainties, is the key to the method's resilience in the presence of sporadic large measurement disturbances. As previously mentioned in section 4.4.3 it is worth emphasizing that, although not shown herein, simultaneous tracking of the spatial periodicity of multiple modes observed at different non-stationary frequencies is possible via adaptive multi-tone tracking in the NLKFFT as exemplified in [Alve 13a]. The proposed method has been benchmarked using JET data from recent campaigns for more than 40 pulses exhibiting notorious MHD activity with results equivalent to the ones presented in figures figure 4.32, figure 4.33 and figure 4.34.

A real-time application for the control of external magnetic perturbations has been developed, implemented and integrated in the suite of real-time controllers in the JET tokamak. The main features of the system are its ease of use from the SL interface, mode-lock based adaptation and slew rate limitation of the current references, multiplexing of references between pre-programmed configuration and externally provided ones in real-time and the choice between PID-based or model-based control.

Although using hardware technology from the 1990s, which eventually became widely adopted and a standard at JET, the system is not only capable of meeting its requirements but also of implementing advanced control schemes within a $500\mu\text{s}$ control cycle period. The use of the MARTe framework for real-time applications greatly eased the development and testing process by enabling development work in a non-real-time environment (Linux) and the deployment in the live environment with minimal code changes. The synchronization mechanism devised enables the application to run at a clock speed twice as fast as the central timing signal available in the VME crate while still providing accurate and consistent time stamping. The KF-based current estimation provides a good improvement ($1/6$ of the raw measurement's noise variance) in terms of SNR of the control variable. This is crucial as all control actions are ultimately driven by it. The mode-lock compensation and slew rate limitation for adjusting in real-time the current references are expected to anticipate and ameliorate events that would, if disregarded, ordinarily stop the entire experiment. In particular, both allow the resume of the originally planned operation conditions in case the offending event is successfully mitigated.

One of the major improvements of this system over its predecessors is the capability of receiving the current references, in real-time, from RTCC. This system allows the implementation of algorithms built upon Simulink-like control/signal processing blocks and used for experimental purposes. This allows the use of arbitrary sets of signals available in RTDN to steer the EFCC currents, in real-time, according to the goals and the evolving state of the plasma pulse. There is already an interest of exploring this feature for the real-time control of the ELM size/frequency crucial to reduce heat loads on to plasma facing components. Finally, the development of an MPC-based control

aimed at the optimization of the controller performance especially in RFA experiments is expected to improve the quality and the overall bandwidth of the system. Although this MPC-based control scheme requires some re-tuning after the move from a cycling frequency of 500 Hz to 2 kHz, this system has already demonstrated to provide the facility to incorporate external magnetic perturbations into plasma experiments at JET.

Although most of the applications presented herein were not actually implemented in real-time, they were designed and tested under the premise of working under real-time conditions.

6.1 Future work

The work presented in this thesis should be considered as a work in progress since a substantial part of it is yet to find the actual real-time hardware/software implementation since it has not been possible to integrate/test in the tight scheduled campaigns of present machines e.g. JET. However, one should stress two possible developments in the short-medium term.

Of all the applications presented in this dissertation, two are notorious candidates for further developments. As previously stated, the controller application of the external magnetic perturbations in the JET tokamak still requires some re-tuning of the VCVS model. In addition, there is a great interest in exploring the real-time experimental control of ELM frequency/size as a proof of principle to ameliorate the heat deposition onto the plasma facing components.

Secondly, the KF implementations for the real-time estimation of the frequency and spatial periodicity of MHD perturbations should be targeted at dedicated devices since it fits perfectly within the prioritization identified on the MHD control highlights of the Eurofusion consortium activities in the fusion roadmap. This stems from the need to acquire sufficient competencies and strategies for the feedback stabilization of MHD activities namely neoclassical tearing modes. Any possible robust strategy for mitigating/controlling MHD instabilities ultimately requires robust and flexible signal processing techniques. It would be of great interest to try these approaches in tokamak devices, such as ASDEX Upgrade and TCV, able to perform real-time feedback control on these instabilities using RF resonance heating and current drive.

Bibliography

- [Agha 05] R. Aghazadeh, H. Lesani, M. Sanaye-Pasand, and B. Ganji. “New technique for frequency and amplitude estimation of power system signals”. *Generation, Transmission and Distribution, IEE Proceedings-*, Vol. 152, No. 3, pp. 435–440, 2005.
- [Alve 04] D. Alves, A. Stephen, N. Hawkes, S. Dalley, A. Goodyear, R. Felton, E. Joffrin, and H. Fernandes. “Real-time motional Stark effect in jet”. *Fusion Engineering and Design*, Vol. 71, No. 14, pp. 175 – 181, 2004. 4th IAEA Technical Meeting on Control, Data Acquisition, and Remote Participation for Fusion Research.
- [Alve 06] D. Alves, R. Coelho, S. R. Cortes, N. C. Hawkes, and J. E. contributors. “MSE diagnostic signal processing using software Phase Locked Loop and Empirical Mode Decomposition”. In: *33rd EPS Conference on Plasma Phys. Rome*, p. 4, 6 2006.
- [Alve 10] D. Alves, R. Coelho, A. Klein, T. Panis, A. Murari, and J. E. Contributors. “A Real-Time Synchronous Detector for the TAE Antenna Diagnostic at JET”. *Nuclear Science, IEEE Transactions on*, Vol. 57, No. 2, pp. 577 –582, april 2010.
- [Alve 11a] D. Alves and R. Coelho. “An Adaptive Algorithm for Real-Time Multi-Tone Estimation and Frequency Tracking of Non-Stationary Signals”. *Nuclear Science, IEEE Transactions on*, Vol. 58, No. 4, pp. 1582 –1587, aug. 2011.
- [Alve 11b] D. Alves, R. Vitelli, L. Zaccarian, L. Zabeo, A. Neto, F. Sartori, P. McCullen, and P. Card. “The new Error Field Correction Coil controller system in the Joint European Torus tokamak”. *Fusion Engineering and Design*, Vol. 86, No. 68, pp. 1034 – 1038, 2011.
- [Alve 12a] D. Alves, R. Felton, S. Jachmich, P. Lomas, P. McCullen, A. Neto, D. F. Valcárcel, G. Arnoux, P. Card, S. Devaux, A. Goodyear, D. Kinna, A. Stephen, and K.-D. Zastrow. “Vessel thermal map real-time system for the JET tokamak”. *Phys. Rev. ST Accel. Beams*, Vol. 15, p. 054701, May 2012.
- [Alve 12b] D. Alves, A. Neto, D. Valcarcel, R. Felton, J. Lopez, A. Barbalace, L. Boncagni, P. Card, G. De Tommasi, A. Goodyear, S. Jachmich, P. Lomas, F. Maviglia, P. McCullen, A. Murari, M. Rainford, C. Reux, F. Rimini, F. Sartori, A. Stephen, J. Vega, R. Vitelli, L. Zabeo, and K.-D. Zastrow. “A new generation of real-time systems in the JET tokamak”. In: *Real Time Conference (RT), 2012 18th IEEE-NPSS*, pp. 1–9, 2012.

- [Alve 13a] D. Alves, R. Coelho, and J.-E. Contributors. “A real-time algorithm for the harmonic estimation and frequency tracking of dominant components in fusion plasma magnetic diagnostics”. *Review of Scientific Instruments*, Vol. 84, No. 8, p. 083508, 2013.
- [Alve 13b] D. Alves, R. Coelho, and the JET EFDA contributors. “Kalman filter methods for real-time frequency and mode number estimation of MHD activity in tokamak plasmas”. *Plasma Physics and Controlled Fusion*, Vol. 55, No. 10, p. 105003, 2013.
- [Ambr 05] G. Ambrosino and R. Albanese. “Magnetic control of plasma current, position, and shape in Tokamaks: a survey or modeling and control approaches”. *Control Systems, IEEE*, Vol. 25, No. 5, pp. 76–92, 2005.
- [Ario 05] M. Ariola and A. Pironti. “Plasma shape control for the JET tokamak: an optimal output regulation approach”. *Control Systems, IEEE*, Vol. 25, No. 5, pp. 65–75, 2005.
- [Arno 12] G. Arnoux, S. Devaux, D. Alves, I. Balboa, C. Balorin, N. Balshaw, M. Beldishevski, P. Carvalho, M. Clever, S. Cramp, J.-L. de Pablos, E. de la Cal, D. Falie, P. Garcia-Sanchez, R. Felton, V. Gervaise, A. Goodyear, A. Horton, S. Jachmich, A. Huber, M. Jouve, D. Kinna, U. Kruezi, A. Manzanares, V. Martin, P. McCullen, V. Moncada, K. Obrejan, K. Patel, P. J. Lomas, A. Neto, F. Rimini, C. Ruset, B. Schweer, G. Sergienko, B. Sieglin, A. Soletto, M. Stamp, A. Stephen, P. D. Thomas, D. F. Valcerel, J. Williams, J. Wilson, K.-D. Zastrow, and J.-E. Contributors. “A protection system for the JET ITER-like wall based on imaging diagnostics”). *Review of Scientific Instruments*, Vol. 83, No. 10, pp. –, 2012.
- [Athe 99] D. Atherton and S. Majhi. “Limitations of PID controllers”. In: *American Control Conference, 1999. Proceedings of the 1999*, pp. 3843–3847 vol.6, 1999.
- [Bara 04] O. Barana, A. Murari, E. Joffrin, F. Sartori, and C. to the EFDA-JET Workprogramme. “Real-time calculation of plasma parameters for feedback control in JET”. *Nuclear Fusion*, Vol. 44, No. 2, p. 335, 2004.
- [Barl 01] I. Barlow, M. Bigi, J. Bird, G. Bonizzoni, R. Buttery, R. Clay, M. D. Benedetti, T. Dobbing, T. Gallagher, G. Gervasini, G. Gittini, T. Hender, S. Hotchin, J. Last, E. Lazzaro, M. Lennholm, S. Peruzzo, V. Riccardo, L. Rossi, A. Santagiustina, D. Starkey, A. Stevens, M. Tabellini, A. Tanga, A. Terrington, M. Vincent, P. Watkins, M. Way, and M. Williams. “The error field correction coils on the JET machine”. *Fusion Engineering and Design*, Vol. 5859, No. 0, pp. 189 – 193, 2001.
- [Baru 12] M. Baruzzo, T. Bolzonella, Y. Liu, G. Manduchi, G. Marchiori, A. Soppelsa, M. Takechi, and F. Villone. “RWM control studies on RFX-mod with a limited set of active coils”. *Nuclear Fusion*, Vol. 52, No. 10, p. 103001, 2012.
- [Begh 05] A. Beghi and A. Cenedese. “Advances in real-time plasma boundary reconstruction: from gaps to snakes”. *Control Systems, IEEE*, Vol. 25, No. 5, pp. 44–64, 2005.

- [Beid 91] H. Beides and G. Heydt. “Dynamic state estimation of power system harmonics using Kalman filter methodology”. *Power Delivery, IEEE Transactions on*, Vol. 6, No. 4, pp. 1663–1670, oct 1991.
- [Beke 89] H. van der Beken, B. Green, C. Steed, J. W. Farthing, P. A. McCullen, and J. How. “Level 1 software at JET: a global tool for physics operation”. In: *Fusion Engineering, 1989. Proceedings., IEEE Thirteenth Symposium on*, pp. 201–204 vol.1, 1989.
- [Bell 10] T. Bellizio, G. De Tommasi, P. McCullen, A. Neto, F. Piccolo, F. Sartori, R. Vitelli, and L. Zabeo. “The Software Architecture of the New Vertical-Stabilization System for the JET Tokamak”. *Plasma Science, IEEE Transactions on*, Vol. 38, No. 9, pp. 2465–2473, sept. 2010.
- [Berr 06] J. Berrino, S. Cirant, F. Gandini, G. Granucci, E. Lazzaro, F. Jannone, P. Smeulders, and G. D’Antona. “Automatic real-time tracking and stabilization of magnetic islands in a Tokamak using temperature fluctuations and ECW power”. *Nuclear Science, IEEE Transactions on*, Vol. 53, No. 3, pp. 1009–1014, 2006.
- [Bitm 86] R. R. Bitmead, A. C. Tsoi, and P. J. Parker. “A Kalman Filtering Approach to Short-Time Fourier Analysis”. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, Vol. ASSP-34, No. 6, p. 1493, 1986.
- [Bitt 00] S. Bittanti and S. Savaresi. “On the parametrization and design of an extended Kalman filter frequency tracker”. *Automatic Control, IEEE Transactions on*, Vol. 45, No. 9, pp. 1718–1724, sep 2000.
- [Bitt 91] S. Bittanti, A. J. Laub, and J. C. Willems. *The Riccati Equation*. Springer-Verlag New York, Inc., 1991.
- [Biza 99] J. Bizarro and A. Figueiredo. “The Wigner distribution as a tool for time-frequency analysis of fusion plasma signals: application to broadband reflectometry data”. *Nuclear Fusion*, Vol. 39, No. 1, p. 61, 1999.
- [Blum 08] J. Blum, C. Boulbe, and B. Faugeras. “Real-time plasma equilibrium reconstruction in a Tokamak”. *Journal of Physics: Conference Series*, Vol. 135, No. 1, p. 012019, 2008.
- [Bonc 11] L. Boncagni, Y. Sadeghi, D. Carnevale, G. Mazzitelli, A. Neto, D. Pucci, F. Sartori, F. Piasco, S. Sinibaldi, V. Vitale, R. Vitelli, L. Zaccarian, S. Monaco, and G. Zamborlini. “First Steps in the FTU Migration Towards a Modular and Distributed Real-Time Control Architecture Based on MARTe”. *Nuclear Science, IEEE Transactions on*, Vol. 58, No. 4, pp. 1778–1783, aug. 2011.
- [Butt 03] R. Buttery, T. Hender, D. Howell, R. L. Haye, O. Sauter, D. Testa, and contributors to the EFDA-JET Workprogramme. “Onset of neoclassical tearing modes on JET”. *Nuclear Fusion*, Vol. 43, No. 2, p. 69, 2003.
- [Butt 04] R. Buttery, T. Hender, D. Howell, R. L. Haye, S. Parris, O. Sauter, C. Windsor, and J.-E. Contributors. “On the form of NTM onset scalings”. *Nuclear Fusion*, Vol. 44, No. 5, p. 678, 2004.

- [Caci 96] M. Caciotta and P. Carbone. “Estimation of non-stationary sinewave amplitude via competitive Kalman filtering”. In: *Instrumentation and Measurement Technology Conference, 1996. IMTC-96. Conference Proceedings. Quality Measurements: The Indispensable Bridge between Theory and Reality., IEEE*, pp. 470–473 vol.1, 1996.
- [Cand 09] J. V. Candy. *Bayesian Signal Processing: Classical, Modern and Particle Filtering Methods (Adaptive and Learning Systems for Signal Processing, Communications and Control Series)*. Wiley-Interscience, Apr. 2009.
- [Carv 11] P. Carvalho, P. Duarte, T. Pereira, R. Coelho, C. Silva, and H. Fernandes. “Real-Time Tomography System at ISTTOK”. *Nuclear Science, IEEE Transactions on*, Vol. 58, No. 4, pp. 1427–1432, aug. 2011.
- [Chap 12] I. Chapman, R. L. Haye, R. Buttery, W. Heidbrink, G. Jackson, C. Muscatello, C. Petty, R. Pinsker, B. Tobias, and F. Turco. “Sawtooth control using electron cyclotron current drive in ITER demonstration plasmas in DIII-D”. *Nuclear Fusion*, Vol. 52, No. 6, p. 063006, 2012.
- [Chen 84] F. F. Chen and M. Lieberman. *Introduction to plasma physics and controlled fusion/Francis F.* Plenum Press, New York, 1984.
- [Chen 85] C. Cheng, L. Chen, and M. Chance. “High-n ideal and resistive shear Alfvén waves in tokamaks”. *Annals of Physics*, Vol. 161, No. 1, pp. 21–47, 1985.
- [Chic 94] J. Chicharo and H. Wang. “Power system harmonic signal estimation and retrieval for active power filter applications”. *Power Electronics, IEEE Transactions on*, Vol. 9, No. 6, pp. 580–586, nov 1994.
- [Chir 91] D. Chiron, T. Bonicelli, M. Huart, M. Garribba, P.-L. Mondino, and P. Noll. “Power supplies for the stabilisation of plasma vertical position: recent upgrades and future development”. In: *Fusion Engineering, 1991. Proceedings., 14th IEEE/NPSS Symposium on*, pp. 513–516 vol.1, 1991.
- [Coel 08] R. Coelho and D. Alves. “Real-time estimation of the poloidal wavenumber of ISTTOK tokamak magnetic fluctuations”. *Rev Sci Instrum*, Vol. 79, Jan 2008.
- [Coel 09a] R. Coelho and D. Alves. “Real-Time Lock-In Amplifier Implementation Using a Kalman Filter for Quasi-Periodic Signal Processing in Fusion Plasma Diagnostics”. *Plasma Science, IEEE Transactions on*, Vol. 37, No. 1, pp. 164–170, jan. 2009.
- [Coel 09b] R. Coelho, D. Alves, N. Hawkes, M. Brix, and J. E. Contributors. “Real-time data processing and magnetic field pitch angle estimation of the JET motional Stark effect diagnostic based on Kalman filtering”. *Review of Scientific Instruments*, Vol. 80, No. 6, p. 063504, 2009.
- [Cool 65] J. Cooley and J. Tukey. “An algorithm for the machine calculation of complex Fourier series”. *Mathematics of computation*, pp. 297–301, 1965.

- [Dash 00] P. Dash, R. K. Jena, G. Panda, and A. Routray. “An extended complex Kalman filter for frequency measurement of distorted signals”. *Instrumentation and Measurement, IEEE Transactions on*, Vol. 49, No. 4, pp. 746–753, 2000.
- [De T 11] G. De Tommasi, D. Alves, T. Bellizio, R. Felton, A. Neto, F. Sartori, R. Vitelli, L. Zabeo, R. Albanese, G. Ambrosino, and P. Lomas. “Real-Time Systems in Tokamak Devices. A Case Study: The JET Tokamak”. *Nuclear Science, IEEE Transactions on*, Vol. 58, No. 4, pp. 1420–1426, aug. 2011.
- [Dorm 13] S. Dormido-Canto, J. Vega, J. Ramirez, A. Murari, R. Moreno, J. Lopez, A. Pereira, and J.-E. Contributors. “Development of an efficient real-time disruption predictor from scratch on JET and implications for ITER”. *Nuclear Fusion*, Vol. 53, No. 11, p. 113001, 2013.
- [Eini 99] G. Einicke and L. White. “Robust extended Kalman filtering”. *Signal Processing, IEEE Transactions on*, Vol. 47, No. 9, pp. 2596–2599, Sep. 1999.
- [Evan 08] T. Evans, M. Fenstermacher, R. Moyer, T. Osborne, J. Watkins, P. Gohil, I. Joseph, M. Schaffer, L. Baylor, M. Bcoulet, J. Boedo, K. Burrell, J. de-Grassie, K. Finken, T. Jernigan, M. Jakubowski, C. Lasnier, M. Lehnen, A. Leonard, J. Lonroth, E. Nardon, V. Parail, O. Schmitz, B. Unterberg, and W. West. “RMP ELM suppression in DIII-D plasmas with ITER similar shapes and collisionalities”. *Nuclear Fusion*, Vol. 48, No. 2, p. 024002, 2008.
- [Faso 95a] A. Fasoli, D. Borba, G. Bosia, D. J. Campbell, J. A. Dobbins, C. Gormezano, J. Jacquinet, P. Lavanchy, J. B. Lister, P. Marmillod, J.-M. Moret, A. Santagiustina, and S. Sharapov. “Direct Measurement of the Damping of Toroidicity-Induced Alfvén Eigenmodes”. *Phys. Rev. Lett.*, Vol. 75, pp. 645–648, Jul 1995.
- [Faso 95b] A. Fasoli, J. Lister, S. Sharapov, S. Ali-Arshad, G. Bosia, D. Borba, D. Campbell, N. Deliyannis, J. Dobbins, C. Gormezano, H. Holties, G. Huysmans, J. Jacquinet, A. Jaun, W. Kerner, P. Lavanchy, J.-M. Moret, L. Porte, A. Santagiustina, and L. Villard. “Overview of Alfvén eigenmode experiments in JET”. *Nuclear Fusion*, Vol. 35, No. 12, p. 1485, 1995.
- [Felt 05] R. Felton, E. Joffrin, A. Murari, L. Zabeo, F. Sartori, F. Piccolo, J. Farthing, T. Budd, S. Dorling, P. McCullen, J. Harling, S. Dalley, A. Goodyear, A. Stephen, P. Card, M. Bright, R. Lucock, E. Jones, S. Griph, C. Hogben, M. Beldishevski, M. Buckley, J. Davis, I. Young, O. Hemming, M. Wheatley, P. Heesterman, G. Lloyd, M. Walters, R. Bridge, H. Leggate, D. Howell, K.-D. Zastrow, C. Giroud, I. Coffey, N. Hawkes, M. Stamp, R. Barnsley, T. Edlington, K. Guenther, C. Gowers, S. Popovichev, A. Huber, C. Ingesson, D. Mazon, D. Moreau, D. Alves, J. Sousa, M. Riva, O. Barana, T. Bolzonella, M. Valisa, P. Innocente, M. Zerbini, K. Bosak, J. Blum, E. Vitale, F. Crisanti, E. de la Luna, and J. Sanchez. “Real-time measurement and control at JET experiment control”. *Fusion Engineering and Design*, Vol. 74, No. 1-4, pp. 561–566, 2005. Proceedings of the 23rd Symposium of Fusion Technology.
- [Felt 99] R. Felton, K. Blackler, S. Dorling, A. Goodyear, O. Hemming, P. Knight, M. Lennholm, F. Milani, F. Sartori, and I. Young. “Real-time plasma control

- at JET using an ATM network”. In: *Real Time Conference, 1999. Santa Fe 1999. 11th IEEE NPSS*, pp. 175–181, 1999.
- [Fern 12] A. Fernandes, R. Pereira, A. Neto, D. Valcarcel, D. Alves, J. Sousa, B. Carvalho, V. Kiptily, B. Syme, P. Blanchard, A. Murari, C. Correia, and C. Varandas. “Real-time processing system for the JET hard X-ray and gamma-ray profile monitor enhancement”. In: *Real Time Conference (RT), 2012 18th IEEE-NPSS*, pp. 1–7, 2012.
- [Ferr 92] J. Ferron and E. Strait. “Real time analysis of tokamak discharge parameters”. *Review of scientific instruments*, Vol. 63, No. 10, pp. 4799–4802, 1992.
- [Figu 04] A. Figueiredo, M. Nave, and E. Contributors. “Timefrequency analysis of non-stationary signals in fusion plasmas using the ChoiWilliams distribution”. *Nuclear Fusion*, Vol. 44, No. 10, p. L17, 2004.
- [Figu 06] A. C. A. Figueiredo and J. ao P. S. Bizarro JET EFDA Contributors. “On limitations of the spectrogram in the representation of the amplitude of non-stationary fusion plasma signals”. *Review of Scientific Instruments*, Vol. 77, No. 10, p. 10F509, 2006.
- [Fina 82] *Final Design Report of Poloidal Radial Field Amplifiers*. Holec, October 1982.
- [Fran 91] G. F. Franklin, J. D. Powell, and A. Emami-Naeini. *Feedback control of dynamic systems*. Vol. 320, Addison-Wesley Reading, MA, 1991.
- [Fu 89] G. Y. Fu and J. W. V. Dam. “Stability of the global Alfvén eigenmode in the presence of fusion alpha particles in an ignited tokamak plasma”. *Physics of Fluids B: Plasma Physics*, Vol. 1, No. 12, pp. 2404–2413, 1989.
- [Gian 04] R. Giannella, N. C. Hawkes, R. Jayakumar, M. Makowski, and L. Zabeo. “Point-to-point analysis of MSE data for plasma diagnostics and control”. *Review of Scientific Instruments*, Vol. 75, No. 10, pp. 4247–4250, 2004.
- [Gned 06] B. V. Gnedenko. *The Theory of Probability: And the Elements of Statistics*. Vol. 132, Amer Mathematical Society, 2006.
- [Gon 11] B. Goncalves, J. Sousa, B. Carvalho, A. Rodrigues, M. Correia, A. Batista, J. Vega, M. Ruiz, J. Lopez, R. Rojo, A. Wallander, N. Utzel, A. Neto, D. Alves, and D. Valcarcel. “Engineering Design of ITER Prototype Fast Plant System Controller”. *Nuclear Science, IEEE Transactions on*, Vol. 58, No. 4, pp. 1439–1446, 2011.
- [Grub 94] P. Gruber and J. Todtli. “Estimation of quasiperiodic signal parameters by means of dynamic signal models”. *Signal Processing, IEEE Transactions on*, Vol. 42, No. 3, pp. 552–562, 1994.
- [Grya 08] M. P. Gryaznevich, T. C. Hender, D. F. Howell, C. D. Challis, H. R. Koslowski, S. Gerasimov, E. Joffrin, Y. Q. Liu, S. Saarelma, and J.-E. contributors. “Experimental studies of stability and beta limit in JET”. *Plasma Physics and Controlled Fusion*, Vol. 50, No. 12, p. 124030, 2008.

- [H Zo 07] H. Zohm and G. Gantenbein and F. Leuterer and A. Manini and M. Maraschek and Q. Yu and the ASDEX Upgrade Team. “Control of MHD instabilities by ECCD: ASDEX Upgrade results and implications for ITER”. *Nuclear Fusion*, Vol. 47, No. 3, p. 228, 2007.
- [Hans 09] J. M. Hanson, B. D. Bono, J. P. Levesque, M. E. Mauel, D. A. Maurer, G. A. Navratil, T. S. Pedersen, D. Shiraki, and R. W. James. “A Kalman filter for feedback control of rotating external kink instabilities in the presence of noise”. *Physics of Plasmas*, Vol. 16, No. 5, p. 056112, 2009.
- [Hast 97] R. Hastie. “Sawtooth Instability in Tokamak Plasmas”. *Astrophysics and Space Science*, Vol. 256, No. 1-2, pp. 177–204, 1997.
- [Hawk 06] N. C. Hawkes, M. Brix, and J. E. Contributors. “Calibration and operational experience with the JET motional Stark effect diagnostic”. *Review of Scientific Instruments*, Vol. 77, No. 10, p. 10E509, 2006.
- [Hawk 99] N. C. Hawkes, K. Blackler, B. Viacoz, C. H. Wilson, J. B. Migozzi, and B. C. Stratton. “Design of the Joint European Torus motional stark effect diagnostic”. *Review of Scientific Instruments*, Vol. 70, No. 1, pp. 894–897, 1999.
- [Haye 06] R. J. L. Haye. “Neoclassical tearing modes and their control”. *Physics of Plasmas*, Vol. 13, No. 5, p. 055501, 2006.
- [Hegn 98] C. C. Hegna. “The physics of neoclassical magnetohydrodynamic tearing modes”. *Physics of Plasmas*, Vol. 5, No. 5, pp. 1767–1774, 1998.
- [Hend 81] H. V. Henderson and S. R. Searle. “On deriving the inverse of a sum of matrices”. *Siam Review*, Vol. 23, No. 1, pp. 53–60, 1981.
- [Hirs 94] S. P. Hirshman, D. K. Lee, F. M. Levinton, S. H. Batha, M. Okabayashi, and R. M. Wieland. “Equilibrium reconstruction of the safety factor profile in tokamaks from motional Stark effect data”. *Physics of Plasmas*, Vol. 1, No. 7, pp. 2277–2290, 1994.
- [Hole 07] M. J. Hole and L. C. Appel. “Fourier decomposition of magnetic perturbations in toroidal plasmas using singular value decomposition”. *Plasma Physics and Controlled Fusion*, Vol. 49, No. 12, p. 1971, 2007.
- [Horn 90] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1990.
- [How 80] J. How *et al.* “The JET pulse termination network”. *Fusion Technology*, Vol. 2, pp. 1421–1426, 1980.
- [Huan 98] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, N. C. Yen, C. C. Tung, and H. H. Liu. “The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis”. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, Vol. 454, No. 1971, pp. 903–995, March 1998.

- [In 06a] Y. In, J. S. Kim, D. H. Edgell, E. J. Strait, D. A. Humphreys, M. L. Walker, G. L. Jackson, M. S. Chu, R. Johnson, R. J. L. Haye, M. Okabayashi, A. M. Garofalo, and H. Reimerdes. “Model-based dynamic resistive wall mode identification and feedback control in the DIII-D tokamak”. *Physics of Plasmas*, Vol. 13, No. 6, p. 062512, 2006.
- [In 06b] Y. In, J. S. Kim, D. H. Edgell, E. J. Strait, D. A. Humphreys, M. L. Walker, G. L. Jackson, M. S. Chu, R. Johnson, R. J. L. Haye, M. Okabayashi, A. M. Garofalo, and H. Reimerdes. “Model-based dynamic resistive wall mode identification and feedback control in the DIII-D tokamak”. *Physics of Plasmas*, Vol. 13, No. 6, p. 062512, 2006.
- [In 06c] Y. In, J. S. Kim, D. H. Edgell, E. J. Strait, D. A. Humphreys, M. L. Walker, G. L. Jackson, M. S. Chu, R. Johnson, R. J. L. Haye, M. Okabayashi, A. M. Garofalo, and H. Reimerdes. “Model-based dynamic resistive wall mode identification and feedback control in the DIII-D tokamak”. *Physics of Plasmas*, Vol. 13, No. 6, p. 062512, 2006.
- [Jaku 07] K. Jakubowska, M. Brix, N. C. Hawkes, E. Rachlew, and J. E. contributors. “MSE diagnostic signal processing using software Phase Locked Loop and Empirical Mode Decomposition”. In: *34th EPS Conference on Plasma Phys. Warsaw*, p. , 7 2007.
- [Joff 03] E. Joffrin, F. Crisanti, R. Felton, X. Litaudon, D. Mazon, D. Moreau, L. Zabeo, R. Albanese, M. Ariola, D. Alves, O. Barana, V. Basiuk, A. Bcoulet, M. Bcoulet, J. Blum, T. Bolzonella, K. Bosak, J. M. Chareau, M. de Baar, E. de la Luna, P. de Vries, P. Dumortier, D. Elbeze, J. Farthing, H. Fernandes, C. Fenzi, R. Giannella, K. Guenther, J. Harling, N. Hawkes, T. C. Hender, D. F. Howell, P. Heesterman, F. Imbeaux, P. Innocente, L. Laborde, G. Lloyd, P. J. Lomas, D. C. McDonald, J. Mailloux, M. Mantsinen, A. Messiaen, A. Murari, J. Ongena, F. Orsitto, V. Pericoli-Ridolfini, M. Riva, J. Sanchez, F. Sartori, O. Sauter, A. C. C. Sips, T. Tala, A. Tuccillo, D. V. Ester, K.-D. Zastrow, M. Zerbini, and contributors to the JET EFDA Programme. “Integrated scenario in JET using real-time profile control”. *Plasma Physics and Controlled Fusion*, Vol. 45, No. 12A, p. A367, 2003.
- [Jouv 11] M. Jouve, C. Balorin, D. Kinna, G. Arnoux, P. Carvalho, S. Devaux, P. Thomas, K. Zastrow, J. Veyret, and J. E. Contributors. “Real-time protection of the ITER-like Wall at JET”. October 2011.
- [Juli 97] S. Julier and J. Uhlmann. “A new extension of the Kalman filter to nonlinear systems”. In: *Int. Symp. Aerospace/Defense Sensing, Simul. and Controls, Orlando, FL*, 1997.
- [Kalm 59] R. Kalman. “On the general theory of control systems”. *IRE Transactions on Automatic Control*, Vol. 4, No. 3, pp. 110–110, 1959.
- [Kalm 60] R. E. Kalman. “A New Approach to Linear Filtering and Prediction Problems”. *Transactions of the ASME—Journal of Basic Engineering*, Vol. 82, No. Series D, pp. 35–45, 1960.

- [Kalm 61] R. E. Kalman and R. S. Bucy. “New results in linear filtering and prediction theory”. *Transactions of the ASME. Series D, Journal of Basic Engineering*, Vol. 83, pp. 95–107, 1961.
- [Kami 07] K. Kamiya, N. Asakura, J. Boedo, T. Eich, G. Federici, M. Fenstermacher, K. Finken, A. Herrmann, J. Terry, A. Kirk, B. Koch, A. Loarte, R. Maingi, R. Maqueda, E. Nardon, N. Oyama, and R. Sartori. “Edge localized modes: recent experimental findings and related issues”. *Plasma Physics and Controlled Fusion*, Vol. 49, No. 7, p. S43, 2007.
- [Kim 99] J. S. Kim, D. H. Edgell, J. M. Greene, E. J. Strait, and M. S. Chance. “MHD mode identification of tokamak plasmas from Mirnov signals”. *Plasma Physics and Controlled Fusion*, Vol. 41, No. 11, p. 1399, 1999.
- [King 98] Q. King, H. Brelén, and J. J. Undertaking. “An experimental automatic control facility at JET”. *JOINT EUROPEAN TORUS-PUBLICATIONS-JET P*, 1998.
- [Klei 08a] A. Klein, H. Carfantan, D. Testa, A. Fasoli, J. Snipes, and J. E. Contributors. “A sparsity-based method for the analysis of magnetic fluctuations in unevenly-spaced Mirnov coils”. *Plasma Physics and Controlled Fusion*, Vol. 50, No. 12, p. 125005, 2008.
- [Klei 08b] A. Klein, D. Testa, J. Snipes, A. Fasoli, H. Carfantan, and J.-E. contributors. “Toroidal mode number analysis of degenerate Alfvén Eigenmodes in the active MHD spectroscopy on JET”. 2008.
- [La S 96] B. La Scala and R. Bitmead. “Design of an extended Kalman filter frequency tracker”. *Signal Processing, IEEE Transactions on*, Vol. 44, No. 3, pp. 739–742, mar 1996.
- [Lefe 01] T. Lefebvre, H. Bruyninckx, J. D. Schutter, T. L. (corresponding, H. Bruyninckx, J. De, and S. D. Mechanical. “Kalman Filters for nonlinear systems: a comparison of performance.”. Tech. Rep., *International Journal of Control*, 2001.
- [Lenn 97] M. Lennholm, D. Campbell, F. Milani, S. Puppini, F. Sartori, and B. Tubbing. “Plasma vertical stabilisation at JET using adaptive gain adjustment”. In: *Fusion Engineering, 1997. 17th IEEE/NPSS Symposium*, pp. 539–542 vol.1, oct 1997.
- [Levi 89] F. M. Levinton, R. J. Fonck, G. M. Gammel, R. Kaita, H. W. Kugel, E. T. Powell, and D. W. Roberts. “Magnetic field pitch-angle measurements in the PBX-M tokamak using the motional Stark effect”. *Phys. Rev. Lett.*, Vol. 63, pp. 2060–2063, Nov 1989.
- [Levi 90] F. M. Levinton, G. M. Gammel, R. Kaita, H. W. Kugel, and D. W. Roberts. “Magnetic field pitch angle diagnostic using the motional Stark effect (invited)”. *Review of Scientific Instruments*, Vol. 61, No. 10, pp. 2914–2919, 1990.

- [Lian 07] Y. Liang, H. R. Koslowski, P. R. Thomas, E. Nardon, B. Alper, P. Andrew, Y. Andrew, G. Arnoux, Y. Baranov, M. Bécoulet, M. Beurskens, T. Biewer, M. Bigi, K. Crombe, E. De La Luna, P. de Vries, W. Fundamenski, S. Gerasimov, C. Giroud, M. P. Gryaznevich, N. Hawkes, S. Hotchin, D. Howell, S. Jachmich, V. Kiptily, L. Moreira, V. Parail, S. D. Pinches, E. Rachlew, and O. Zimmermann. “Active Control of Type-I Edge-Localized Modes with $n = 1$ Perturbation Fields in the JET Tokamak”. *Phys. Rev. Lett.*, Vol. 98, p. 265004, Jun 2007.
- [Liu 00] Y. Q. Liu and A. Bondeson. “Active Feedback Stabilization of Toroidal External Modes in Tokamaks”. *Phys. Rev. Lett.*, Vol. 84, pp. 907–910, Jan 2000.
- [Lomb 76] N. Lomb. “Least-squares frequency analysis of unequally spaced data”. *Astrophysics and Space Science*, Vol. 39, pp. 447–462, 1976.
- [Ma 96] H. Ma and A. Girgis. “Identification and tracking of harmonic sources in a power system using a Kalman filter”. *Power Delivery, IEEE Transactions on*, Vol. 11, No. 3, pp. 1659–1665, 1996.
- [Mall 99] S. Mallat. *A wavelet tour of signal processing*. Academic press, 1999.
- [Mand 94] A. Mandelis. “Signal-to-noise ratio in lock-in amplifier synchronous detection: A generalized communications systems approach with applications to frequency, time, and hybrid (rate window) photothermal measurements”. *Review of Scientific Instruments*, Vol. 65, No. 11, pp. 3309–3323, 1994.
- [Masi 02] A. Masiello, M. Bigi, R. Buttery, and V. Riccardo. “Finite element electromagnetic analysis of the jet error field correction coils”. *Fusion Engineering and Design*, Vol. 63-64, No. , pp. 455 – 460, 2002.
- [Mazo 10] D. Mazon, J. Blum, C. Boulbe, B. Faugeras, A. Boboc, M. Brix, P. De-Vries, S. Sharapov, and L. Zabeo. “Equinox: a real-time equilibrium code and its validation at JET”. *From physics to control through an emergent view*, Vol. 15, pp. 327–333, 2010.
- [Mich 41] W. C. Michels and N. L. Curtis. “A Pentode LockIn Amplifier of High Frequency Selectivity”. *Review of Scientific Instruments*, Vol. 12, No. 9, pp. 444–447, 1941.
- [Mora 99] M. Morari and J. H. Lee. “Model predictive control: past, present and future”. *Computers & Chemical Engineering*, Vol. 23, No. 45, pp. 667 – 682, 1999.
- [Nard 92] C. Nardone. “Multichannel fluctuation data analysis by the singular value decomposition method. Application to MHD modes in JET”. *Plasma Physics and Controlled Fusion*, Vol. 34, No. 9, p. 1447, 1992.
- [Neto 10] A. Neto, F. Sartori, F. Piccolo, R. Vitelli, G. De Tommasi, L. Zabeo, A. Barbalace, H. Fernandes, D. Valcarcel, and A. J. N. Batista. “MARTE: A Multiplatform Real-Time Framework”. *Nuclear Science, IEEE Transactions on*, Vol. 57, No. 2, pp. 479–486, 2010.

- [Neto 11] A. Neto, D. Alves, L. Boncagni, P. Carvalho, D. Valcarcel, A. Barbalace, G. De Tommasi, H. Fernandes, F. Sartori, E. Vitale, R. Vitelli, and L. Zabeo. “A Survey of Recent MARTe Based Systems”. *Nuclear Science, IEEE Transactions on*, Vol. 58, No. 4, pp. 1482–1489, 2011.
- [Noll 91] D. Noll, D. Nishimura, and A. Macovski. “Homodyne detection in magnetic resonance imaging”. *Medical Imaging, IEEE Transactions on*, Vol. 10, No. 2, pp. 154–163, 1991.
- [Nune 99] F. D. Nunes, J. Santos, and M. E. Manso. “Recursive algorithm for fast evaluation of the Abel inversion integral in broadband reflectometry”. pp. 1047–1050, AIP, 1999.
- [Orto 06] S. Ortolani and the RFX team. “Active MHD control experiments in RFX-mod”. *Plasma Physics and Controlled Fusion*, Vol. 48, No. 12B, p. B371, 2006.
- [Ou 07] Y. Ou, M. Walker, E. Schuster, and J. Ferron. “Equilibrium reconstruction improvement via Kalman-filter-based vessel current estimation at DIII-D”. *Fusion Engineering and Design*, Vol. 82, No. 5-14, pp. 1144 – 1152, 2007. Proceedings of the 24th Symposium on Fusion Technology - SOFT-24.
- [Poli 08] F. M. Poli, S. E. Sharapov, S. C. Chapman, and J.-E. Contributors. “Study of the spectral properties of ELM precursors by means of wavelets”. *Plasma Physics and Controlled Fusion*, Vol. 50, No. 9, p. 095009, 2008.
- [Rabi 75] L. Rabiner and B. Gold. “Theory and application of digital signal processing”. *Englewood Cliffs, NJ, Prentice-Hall, Inc., 1975. 777 p.*, Vol. 1, 1975.
- [Raju 00] D. Raju, R. Jha, P. Kaw, S. Mattoo, Y. Saxena, and A. Team. “Mirnov coil data analysis for tokamak ADITYA”. *Pramana*, Vol. 55, pp. 727–732, 2000.
- [Roma 13] F. Romanelli and J. E. Contributors. “Overview of the JET results with the ITER-like wall”. *Nuclear Fusion*, Vol. 53, No. 10, p. 104002, 2013.
- [Rout 02] A. Routray, A. Pradhan, and K. Rao. “A novel Kalman filter for frequency estimation of distorted signals in power systems”. *Instrumentation and Measurement, IEEE Transactions on*, Vol. 51, No. 3, pp. 469–479, 2002.
- [S Bo 07] S. Bourguignon, H. Carfantan, and T. Böhm. “SparSpec: a new method for fitting multiple sinusoids with irregularly sampled data”. *A&A*, Vol. 462, No. 1, pp. 379–387, 2007.
- [Saiz 97] V. Saiz and J. Guadalupe. “Application of Kalman filtering for continuous real-time tracking of power system harmonics”. *Generation, Transmission and Distribution, IEE Proceedings-*, Vol. 144, No. 1, pp. 13 –20, jan 1997.
- [Saut 02] O. Sauter, E. Westerhof, M. L. Mayoral, B. Alper, P. A. Belo, R. J. Buttery, A. Gondhalekar, T. Hellsten, T. C. Hender, D. F. Howell, T. Johnson, P. Lamalle, M. J. Mantsinen, F. Milani, M. F. F. Nave, F. Nguyen, A. L. Pecquet, S. D. Pinches, S. Podda, and J. Rapp. “Control of Neoclassical Tearing Modes by Sawtooth Control”. *Phys. Rev. Lett.*, Vol. 88, p. 105001, Feb 2002.

- [Saut 04] O. Sauter. “On the contribution of local current density to neoclassical tearing mode stabilization”. *Physics of Plasmas*, Vol. 11, No. 10, pp. 4808–4813, 2004.
- [Scar 82] J. Scargle. “Studies in astronomical time series analysis. II-Statistical aspects of spectral analysis of unevenly spaced data”. *The Astrophysical Journal*, Vol. 263, pp. 835–853, 1982.
- [Shaf 63] V. D. Shafranov. “Equilibrium of a toroidal plasma in a magnetic field”. *Journal of Nuclear Energy. Part C, Plasma Physics, Accelerators, Thermonuclear Research*, Vol. 5, No. 4, p. 251, 1963.
- [Shi 06] Y. Shi. “Digital lock-in technique for the motional Stark effect diagnostic”. *Review of Scientific Instruments*, Vol. 77, No. 3, p. 036111, 2006.
- [Simo 06] D. Simon. *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches*. Wiley-Interscience, 2006.
- [Snip 04] J. A. Snipes, D. Schmittiel, A. Fasoli, R. S. Granetz, and R. R. Parker. “Initial active MHD spectroscopy experiments exciting stable Alfvén eigenmodes in Alcator C-Mod”. *Plasma Physics and Controlled Fusion*, Vol. 46, No. 4, p. 611, 2004.
- [Sonn 64] L. Sonneborn and F. Van Vleck. “The bang-bang principle for linear control systems”. *Journal of the Society for Industrial & Applied Mathematics, Series A: Control*, Vol. 2, No. 2, pp. 151–159, 1964.
- [Step 11] A. Stephen, G. Arnoux, T. Budd, P. Card, R. Felton, A. Goodyear, J. Harling, D. Kinna, P. Lomas, P. McCullen, and J. E. contributors. “Centralised Coordinated Control To Protect The JET ITER-like Wall”. p. , October 2011.
- [Stra 99] B. C. Stratton, D. Long, R. Palladino, and N. C. Hawkes. “Instrumentation for the joint European torus motional Stark effect diagnostic”. *Review of Scientific Instruments*, Vol. 70, No. 1, pp. 898–901, 1999.
- [Test 04] D. Testa, A. Fasoli, D. Borba, A. Jaun, S. Sharapov, and J. E. Contributors. “Experimental studies of Alfvén mode stability in the JET tokamak”. In: *Proc. 23rd Symp. Fusion Technol. (SOFT), Venice, Italy, 9 2004*.
- [Test 07] D. Testa, A. Fasoli, A. Klein, P. Beaumont, R. Bertizzolo, R. Chavan, F. Dolizy, P. Lavanchy, P. Marmillod, V. Riccardo, S. Sanders, J. Snipes, J. Thomas, P. Titus, R. Walton, and M. W. J. E. contributors. “First measurement of the damping rate of high-n toroidal Alfvén eigenmodes in JET tokamak plasma”. 2007.
- [Test 10] D. Testa, H. Carfantan, A. Goodyear, P. Blanchard, A. Klein, T. Panis, *et al.* “The JET Alfvén Eigenmode Local Manager for the real-time detection and tracking of MHD instabilities”. *EPL (Europhysics Letters)*, Vol. 92, No. 5, p. 50001, 2010.
- [Tref 97] L. Trefethen and D. Bau III. *Numerical linear algebra*. Society for Industrial Mathematics, 1997.

- [Valc 11] D. Valcarcel, A. Neto, I. Carvalho, B. Carvalho, H. Fernandes, J. Sousa, F. Janky, J. Havlicek, R. Beno, J. Horacek, M. Hron, and R. Panek. “The COMPASS Tokamak Plasma Control Software Performance”. *Nuclear Science, IEEE Transactions on*, Vol. 58, No. 4, pp. 1490–1496, aug. 2011.
- [Valc 12] D. Valcarcel, D. Alves, A. Neto, C. Reux, B. Carvalho, R. Felton, P. Lomas, J. Sousa, and L. Zabeo. “Parallel task management library for MARTe”. In: *Real Time Conference (RT), 2012 18th IEEE-NPSS*, pp. 1–7, 2012.
- [Wan 00] E. Wan and R. Van Der Merwe. “The unscented Kalman filter for nonlinear estimation”. In: *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, pp. 153–158, 2000.
- [Wess 11] J. Wesson. *Tokamaks*. Vol. 149, OUP Oxford, 2011.
- [Woma 83] K. H. Womack. “Interferometric Phase Measurement Using Spatial Synchronous Detection”. pp. 8–15, 1983.
- [Wong 99] K.-L. Wong. “A review of Alfvén eigenmode observations in toroidal plasmas”. *Plasma Physics and Controlled Fusion*, Vol. 41, No. 1, p. R1, 1999.
- [Xion 08] K. Xiong, H. Zhang, and L. Liu. “Adaptive robust extended Kalman filter for nonlinear stochastic systems”. *Control Theory Applications, IET*, Vol. 2, No. 3, pp. 239–250, March 2008.
- [Zabe 02] L. Zabeo, A. Murari, E. Joffrin, D. Mazon, and C. Taliencio. “A versatile method for the real time determination of the safety factor and density profiles in JET”. *Plasma Physics and Controlled Fusion*, Vol. 44, No. 11, p. 2483, 2002.
- [Zakh 73] L. E. Zakharov and V. D. Shafranov. “Equilibrium of a toroidal plasma with noncircular cross section”. *Soviet Physics Technical Physics*, Vol. 18, p. 151, Aug. 1973.
- [Zege 06] S. Zegenhagen, A. Werner, A. Weller, and T. Klinger. “Analysis of Alfvén eigenmodes in stellarators using non-evenly spaced probes”. *Plasma Physics and Controlled Fusion*, Vol. 48, No. 9, p. 1333, 2006.
- [Zohm 07] H. Zohm, G. Gantenbein, F. Leuterer, A. Manini, M. Maraschek, Q. Yu, and the ASDEX Upgrade Team. “Control of MHD instabilities by ECCD: ASDEX Upgrade results and implications for ITER”. *Nuclear Fusion*, Vol. 47, No. 3, p. 228, 2007.