New particle filters for underwater terrain-aided navigation using multi-sensor fusion

Camille Palmier

université de BORDEAUX

THÈSE PRÉSENTÉE

POUR OBTENIR LE GRADE DE

DOCTEUR
DE L'UNIVERSITÉ DE BORDEAUX

ECOLE DOCTORALE MATHÉMATIQUES ET INFORMATIQUE

SPÉCIALITÉ MATHÉMATIQUES APPLIQUÉES

Par **Camille PALMIER**

New particle filters for underwater terrain-aided navigation using multi-sensor fusion

Sous la direction de : **Pierre DEL MORAL**
Co-encadrants : **Karim DAHIA et Dann LANEUVILLE**

Membres du jury :

Mme. Audrey GIREMUS	Professeure	Université de Bordeaux	Présidente
M. François LE GLAND	Directeur de recherche	Inria Rennes	Rapporteur
M. Fredrik GUSTAFSSON	Professeur	Université de Linköping	Rapporteur
Mme. Lyudmila MIHAYLOVA	Professeure	Université de Sheffield	Examinatrice
M. Pierre DEL MORAL	Directeur de recherche	Inria Bordeaux	Directeur
M. Karim DAHIA	Ingénieur de recherche	ONERA Palaiseau	Encadrant
M. Dann LANEUVILLE	Ingénieur d'étude	Naval Group Nantes	Encadrant invité
M. Nicolas MERLINGE	Ingénieur de recherche	ONERA Palaiseau	Invité

ACKNOWLEDGMENTS

J'aimerais remercier toutes les personnes qui m'ont aidées durant ces trois dernières années et ont grandement contribuées à la réussite de cette thèse.

Je souhaiterais remercier François Le Gland et Fredrik Gustafsson pour avoir accepté d'être rapporteurs de cette thèse. Vos critiques constructives ont permis d'améliorer ce manuscrit. Merci également à Lyudmila Mihaylova d'avoir accepté d'examiner cette thèse ainsi que pour les questions très intéressantes posées lors de la soutenance, et une nouvelle fois merci à Audrey Giremus d'avoir présidé le jury et de m'avoir permis de mettre en évidence des pistes de recherche prometteuses.

Je tiens à remercier Laurent Michel et Emmanuel Gobet pour avoir accepté de faire partie de mon comité de suivi individuel. Nos réunions m'ont permis d'avancer efficacement dans mes travaux de recherche et vos conseils notamment sur l'organisation de la fin de thèse m'ont été très précieux.

J'aimerais maintenant remercier mes encadrants. Merci à Pierre Del Moral pour avoir accepté de diriger cette thèse et pour m'avoir intégrée dans des travaux de recherche théorique prestigieux. Je souhaiterais également remercier Dann Laneuville pour son aide sur les aspects appliqués, pour son enthousiasme et ses encouragements. Un remerciement spécial à Karim Dahia pour son immense soutien tout au long de la thèse. Je tiens également à remercier ici Nicolas Merlinge qui, bien que ne faisant pas partie de mon encadrement officiel, a toujours répondu présent pour m'aider à avancer. Merci à vous deux pour m'avoir soutenu sur le plan professionnel comme personnel pendant ces trois ans !

J'aimerais remercier les équipes que j'ai eu l'occasion de côtoyer à l'ONERA et à Naval Group. Merci à l'équipe de Naval Group qui m'a accueillie à l'occasion de quelques jours de team building. Je garde de très bon souvenirs de ces moments ! Je tiens à remercier Gilbert Brutus Hourié pour être venu plusieurs fois à Nantes afin d'échanger sur mes travaux et leur dimension appliquée. Merci à l'unité CEVA pour m'avoir accueillie durant ces trois ans. Un remerciement particulier à Christian Musso et à Jean-Michel Allard pour avoir suivi mes travaux, proposé des pistes de recherche, validé des points techniques et fait de nombreuses relectures. Un grand merci à Hervé Verrièle sans qui la thèse n'aurait jamais vu le jour.

Ces dernières années ont été particulièrement riches en rencontres et j'aimerais remercier tous les doctorants et les stagiaires de l'ONERA pour avoir assuré une très bonne ambiance au quotidien. Merci d'abord aux "anciens" : Vincent, Iréna, Nathan, Sergio, Émilien et Camille. Merci Émilien pour tes conseils ! J'ai adoré être ton binôme infernal et répandre la terreur lors des désignations petit-déj. Merci Camille pour m'avoir motivée et remontée le moral dans les moments difficiles ! Merci aussi pour m'avoir initiée à la boxe et m'avoir vengée lors de combats trop déséquilibrés. Merci aux moins anciens (mais quand même) : Julien (Bulien et Mulien) et Ali. Merci Julien pour tous les moments que nous avons passés ensemble et désolée pour les litres de thé renversés dans ton bureau. Amiral, ne t'inquiète

pas je sais que tu as toujours raison ! Un grand merci Ali pour ta gentillesse, tes conseils lecture et tes entraînements de foot. N'oublie pas ta promesse ! Un grand merci à ceux qui ont été là du début à la fin : Enzo, Esteban, Guillaume, Javiera et Baptiste. Merci Enzo (le "meilleur co-bureau") pour m'avoir aidée à améliorer mes réflexes et mes compétences informatiques. Merci pour ta patience infinie lorsque j'essaie de faire des figures gnuplot "seule" ou de me connecter "seule" à "mon" serveur. J'espère venir tous les ans skier avec toi à Cham. Merci Esteban (alias viking-vampire-enfant-du-soleil ou vrai-faux ami) d'avoir aimé mes musiques et de les avoir jouées à la guitare. J'espère qu'elles ne te hanteront pas pour toujours ! Je garde de très bons souvenirs de nos soirées films et des voyages que nous avons partagés : "you're like cinnamon". Merci aux "jeunes" : Julius (super soirées karaoké), Antonello, Étienne, Clément (très bons souvenirs de notre voyage à Sun City), Hanae (très bonne co-bureau et future "maman" des doctorants), Clara, Mathieu, Juliette, Paul, Baptiste, Jorge, Renato, Romain, Thomas, Pelin, Jules, Périclès (j'espère n'oublier personne). Merci Dr "stagiaire" Thomas pour avoir égayé ma fin de thèse, d'avoir été mon confident et pour toute ton aide.

Je souhaiterais remercier deux amis Raphaël et Olivier. Merci Raphaël de m'avoir écouté et motivé durant ces trois ans. Merci Olivier pour toute l'aide que tu m'as apportée. Merci d'avoir passé de nombreuses heures à écouter mes journées et à relire mes articles et mon manuscrit !

Je tiens à remercier ceux qui m'ont mis sur la voie de la recherche : parrain Philippe et Julien Worms. Merci parrain pour m'avoir transmis la passion des sciences et de la recherche ! Merci Julien pour avoir cru en moi, m'avoir aidé dans mes choix d'orientation et pour continuer de prendre des nouvelles.

Finalement, mes derniers remerciements vont à ma famille sans qui rien ne serait possible et qui ont été une source de motivation au quotidien. Merci pour tout !

Nouveaux filtres particulaires pour la navigation sous-marine par fusion multi-capteurs

Résumé : Les travaux présentés dans ce mémoire de thèse portent sur le développement et l'étude de filtres particulaires robustes pour la navigation sous-marine par corrélation de terrain. En effet, les filtres développés permettent de contrôler l'erreur de l'approximation Monte Carlo due aux approximations d'intégrales et à l'étape de ré-échantillonnage.

Une première stratégie consiste à maintenir la cohérence entre la vraisemblance et la densité prédite en approchant la vraisemblance par un noyau adaptatif. Ainsi, les poids des particules dégénèrent plus lentement ce qui permet de réduire les cas de divergence du filtre. Cette méthode est appelé Adaptive Approximate Bayesian Computation (A2BC) et a été intégrée aux filtres Regularised Particle Filter (RPF) et Rao-Blackwellised Particle Filter (RBPF).

Une seconde stratégie est fondée sur le choix d'une densité d'importance dont le support recouvre celui de la densité conditionnelle. Cette stratégie a été intégrée à un filtre particulaire qui se nomme l'Interacting Weighted Ensemble Kalman Filter (IWEnKF) et qui a pour base le Weighted Ensemble Kalman Filter (WEnKF). L'IWEnKF calcule analytiquement le support de la densité d'importance afin d'assurer un recouvrement optimal avec la densité conditionnelle, ce qui permet de réduire les fluctuations Monte Carlo.

Ces nouveaux filtres ont été appliqués au recalage de navigation inertielle d'un véhicule sous-marin par corrélation de terrain. Le véhicule sous-marin est équipé d'un sondeur multi-faisceaux, d'un gravimètre et de cartes numériques embarquées associées aux capteurs stockées dans le calculateur de bord. Les résultats obtenus montrent une nette amélioration en termes de précision et de robustesses pour les filtres A2BC et l'IWEnKF par rapport aux filtres particulaires classiques (RPF, RBPF et WEnKF).
Mots clés : Navigation sous-marine, Filtres particulaires, Fusion multi-capteur, Corrélation de terrain

---

New particle filters for underwater terrain-aided navigation using multi-sensor fusion

Abstract: The goal of this thesis is to develop and study robust particle filters for underwater terrain aided navigation. The studied filters allow the control of the Monte Carlo approximation errors that are due to the evaluation of the integrals and to the resampling step.

The first strategy consists in maintaining the consistency between the likelihood and the prior density by adapting the support of the likelihood. In that way, the particles' weights degeneracy is slowed down, which reduces the resampling frequency and thus the cases of divergence. This approach is called Adaptive Approximate Bayesian Computation (A2BC) and is integrated within the Regularised Particle Filter (RPF) and the Rao-Blackwellised Particle Filter (RBPF).

The second strategy is based on the choice of the importance density whose support overlaps the conditional density. The proposed filter is called the Interacting Weighted Ensemble Kalman Filter (IWEnKF), and is based on the Weighted Ensemble Kalman Filter (WEnKF). IWEnKF analytically computes the support of the importance density to ensure an optimal overlap with the conditional density, therefore reducing Monte Carlo fluctuations.

The proposed filters (A2BC-Particle Filters and IWEnKF) were applied to an underwater terrain-aided navigation case. The underwater vehicle was equipped with a multi-beam telemeter, an atomic gravimeter, and numerical maps associated with the sensors stored in an on-board computing system. The results show an improvement of the accuracy and of the robustness to non-linearities for the A2BC-Particle Filters and the IWEnKF compared to conventional particle filters (RPF, RBPF, and WEnKF).
Keywords: Underwater navigation, Particle filters, Multi-sensor fusion, Terrain-aided navigation

---

Institut de Mathématiques de Bordeaux
UMR 5251, Université de Bordeaux, F-33400 Talence, France

RÉSUMÉ ÉTENDU EN FRANÇAIS

Le problème d'estimation consiste à restituer l'état courant inconnu d'un système à partir de mesures bruitées et partielles. Des incertitudes sont introduites dans le processus d'estimation à cause de la méconnaissance des modèles de dynamique et d'observation. Le problème d'estimation est résolu par un filtre qui fournit à chaque instant une estimation de l'état. Dans le cas d'une application embarquée, le temps d'exécution de ce filtre doit être acceptable. Dans ce contexte, les filtres récursifs sont ainsi privilégiés.

Dans le cas linéaire gaussien, la densité de l'état reste gaussienne et la solution optimale est donnée par le filtre de Kalman. Ce dernier approche la densité de l'état sachant les mesures (densité conditionnelle) par une densité gaussienne caractérisée par sa moyenne et sa matrice de covariance. La non-linéarité du modèle peut entraîner la multi-modalité de la densité conditionnelle et ainsi mettre en échec le filtre de Kalman. Lorsque le système est fortement non-linéaire, le filtre de Kalman étendu peut également diverger.

Au début des années 90, des méthodes numériques ont été proposées sous le nom de filtres particulaires. Le filtrage particulaire permet de traiter avec succès des problèmes d'estimation non-linéaires dans lesquels l'incertitude des mesures est soumise à des lois statistiques quelconques. Le filtrage particulaire permet d'approcher la densité conditionnelle par une somme pondérée de Dirac. Un ensemble de points appelés particules est simulé; chaque particule représente un état probable du système. Les coefficients de pondération (poids) associés à chaque particule sont une mesure du degré de confiance que l'on peut avoir en ces dernières pour représenter l'état réel. A chaque itération, les particules évoluent suivant l'équation d'état du système (étape de prédiction) et les poids sont mis à jour en fonction des mesures (étape de correction).

Cependant, la convergence uniforme du filtre particulaire classique n'est malheureusement pas assurée. En effet, le choix de la densité de transition pour déplacer les particules n'est pas toujours optimal ce qui a pour conséquence d'éloigner les particules de la zone d'intérêt. De ce fait, au bout de quelques itérations, les poids des particules dégénèrent. Cela signifie que les poids décroissent vers 0 pour toutes les trajectoires simulées à l'exception d'une qui a un poids proche de 1. Afin de remédier à ce problème, une étape de ré-échantillonnage des particules est effectuée. Cette étape consiste à générer un autre système de particules de même taille en favorisant les particules ayant un poids important. Les particules ayant un poids important sont dupliquées et les autres sont éliminées.

Certaines améliorations du filtre particulaire ont été proposées comme le Regularised Particle Filter (RPF) et le Rao-Blackwellised Particle Filter (RBPF) afin respectivement d'améliorer la robustesse dans le cas d'un faible bruit de dynamique et de gérer les grandes dimensions de l'espace d'état.

Malgré ces améliorations, les filtres particulaires peuvent diverger à cause des approximations de Monte Carlo réalisées successivement lors de l'évaluation des intégrales et lors de

l'étape de ré-échantillonnage. Le bon fonctionnement du filtre particulaire dépend fortement du degré de cohérence entre la densité prédite et la vraisemblance d'une part et du choix de la densité d'importance pour déplacer les particules vers la zone d'intérêt d'autre part. Ainsi des choix appropriés de vraisemblance et de densité d'importance permettent de concevoir des filtres plus robustes aux erreurs de capteurs. Afin de réduire l'erreur de Monte Carlo (la variance des filtres), nous proposons deux approches :

- approcher la vraisemblance par une densité lisse (noyau) à support adaptatif qui permet d'assurer la cohérence avec le support de la densité prédite. A chaque pas de temps, la cohérence entre la densité prédite et le noyau est contrôlée par un paramètre qui permet d'ajuster la fenêtre du noyau. La méthode que nous proposons, nommée Adaptive Approximate Bayesian Computation (A2BC), utilise l'indicateur de ré-échantillonnage afin de calculer ce paramètre. Cela permet d'assurer la cohérence entre la densité prédite et le noyau sans dégrader l'information issue de la mesure. La méthode A2BC permet de garder les particules dans les différents modes (maxima) de la vraisemblance le plus longtemps possible en attendant de parcourir un relief plus favorable au recalage afin d'éviter les cas de divergence. Les poids des particules dégénèrent donc plus lentement. Le nombre nécessaire de ré-échantillonnages du système de particules est alors réduit. L'A2BC est fondée sur la méthode Approximate Bayesian Computation (ABC) qui permet de résoudre le problème d'estimation lorsque la vraisemblance est inconnue ou qu'elle est trop coûteuse à évaluer. L'A2BC est une méthode qui peut être intégrée dans n'importe quel filtre particulaire. Dans ces travaux, elle a été intégrée dans le RPF et le RBPF. Les résultats théoriques montrent que :
  - l'erreur locale de Monte Carlo est bornée par le critère de cohérence entre la densité prédite et le noyau qui est contrôlé par la méthode A2BC;
  - la variance des poids non normalisés admet une borne supérieure qui est directement contrôlée par la méthode A2BC.
- choisir une densité d'importance qui garantisse un meilleur recouvrement avec le support de la densité conditionnelle. Le choix de la densité d'importance est crucial pour le bon fonctionnement des filtres particulaires. L'algorithme d'estimation proposé, nommé Interacting Weighted Ensemble Kalman Filter (IWEnKF), est fondé sur le filtre de Kalman d'Ensemble pondéré (WEnKF). Le WEnKF peut être considéré comme un filtre particulaire avec une densité d'importance gaussienne issue d'un filtre de Kalman. Afin de garantir un recouvrement optimal entre le support de la densité d'importance et le support de la densité conditionnelle et ainsi de réduire l'erreur de Monte Carlo, nous proposons une approche analytique utilisant une décomposition polaire pour calculer le support de la densité d'importance. De cette façon, l'IWEnKF est plus robuste aux non-linéarités que le WEnKF.

L'application considérée dans ces travaux de thèse est la navigation sous-marine. La capacité d'un véhicule sous-marin autonome à accomplir une mission dépend des performances des algorithmes de navigation embarqués. La navigation de ces véhicules utilise la plupart du temps des mesures issues d'une centrale inertielle. Bien que la navigation inertielle soit autonome et fiable, elle produit des mesures imparfaites (par exemple sujettes à des biais et à des bruits) qui mènent à une dérive croissante avec le temps dans la solution de navigation. Ainsi, si la navigation inertielle suffit généralement à la réalisation d'une mission de courte durée, des données exogènes de recalage sont nécessaires pour une mission longue. Les mesures GPS sont indisponibles à cause de l'environnement sous-marin et refaire surface est souvent exclu pour des questions de discrétion.

Une méthode de recalage couramment utilisée est la navigation par corrélation de terrain. Cette méthode permet d'estimer l'état d'un véhicule sous-marin (par exemple sa position, sa vitesse et son attitude) en corrélant un profil du terrain reconstruit à partir des mesures d'un capteur avec un profil d'une carte embarquée de la zone d'opérations. Le filtrage particulaire s'accommode de ce contexte. En effet :

- les équations de mesures qui relient à chaque instant les mesures bathymétrique et gravimétrique à l'état du système sont fortement non-linéaires;
- plusieurs trajectoires d'un véhicule sous-marin peuvent avoir les mêmes profils de profondeur ou de gravité pendant un certain nombre de mesures étant donnée l'ambiguïté des cartes de terrain. Cela se traduit par la multi-modalité de la densité conditionnelle et est la raison de l'échec du filtre de Kalman étendu;
- la dimension de l'espace d'état à estimer est grande ce qui exclut l'utilisation des méthodes de maillage classiquement utilisées.

Le capteur utilisé pour la méthode de corrélation de terrain est un sondeur multi-faisceaux qui produit une série de mesures de profondeur au cours de la trajectoire du véhicule. Le capteur multi-faisceaux est couplé à une carte de profondeur. Si le terrain contient suffisamment d'informations, ce type de capteur permet de retrouver la position. L'absence de mesure directe de la vitesse du véhicule ne permet pas d'en fournir une estimation précise, car elle est obtenue par intégration des équations du modèle. Pour améliorer l'estimation des paramètres cinématiques du véhicule, un autre capteur, prometteur pour la navigation sous-marine, est introduit : un gravimètre à atomes froids. Le gravimètre développé à l'ONERA permet en particulier d'avoir des mesures de gravité précises. Le gravimètre à atomes froids est couplé à une carte d'anomalies de gravité. L'ajout du gravimètre permet d'améliorer les estimations de la position et de la vitesse du véhicule par mesure directe de cette dernière.

La disponibilité des cartes des fonds marins (profondeurs) et des cartes d'anomalies de gravité ainsi que l'apparition de sonars dédiés et de gravimètres à atomes froids précis rendent pertinente l'utilisation de la méthode de corrélation de terrain pour le domaine sous-marin.

Dans le cadre de ces travaux de thèse, les filtres A2BC ont été appliqués au problème de recalage de navigation inertielle. Les performances des filtres particulaires A2BC-RPF et A2BC-RBPF ont été étudiées et comparées à celles des filtres particulaires classiques (RPF et RBPF) pour un nombre de particules identique. La précision et la robustesse des filtres sont améliorées par la méthode A2BC. En effet, les taux de non-convergence des filtres RPF et RBPF (8% et 10% respectivement) sont plus élevés que ceux des filtres A2BC-RPF et A2BC-RBPF (3% et 6% respectivement).

Dans le cas d'un autre scenario avec une fusion centralisée des mesures gravimètrique et bathymétrique, nous montrons une nette amélioration de la précision et de la robustesse par rapport à l'utilisation seule de la bathymétrie. En effet, dans le cas de la fusion, les erreurs d'estimation des positions et des vitesses convergent plus rapidement et vers des valeurs plus faibles que dans le cas de la bathymétrie seule. Au niveau de la robustesse, le RPF affiche un taux de non-convergence de 12% dans le cas de la bathymétrie seule et de 7% dans le cas de la fusion. Le pourcentage de non-convergence est encore réduit dans le cas des filtres particulaires A2BC. Le A2BC-RPF a un taux de non-convergence de 8% dans le cas de la bathymétrie seule et de 3% dans le cas de la fusion.

Les perfomances du filtre IWEnKF ont été comparées à celle du WEnKF sur un problème de recalage de navigation par mesures bathymétriques pour un nombre de particules identique. Les résultats numériques montrent une meilleure robustesse du filtre IWEnKF par rapport au WEnKF, le taux de non-convergence du WEnKF étant de 12% contre 6% pour le IWEnKF.

Les algorithmes proposés (A2BC-PF et IWEnKF) peuvent s'adapter à de nombreux problèmes de filtrage non-linéaire. La simplicité algorithmique de la méthode A2BC lui permet d'être intégrée dans n'importe quel filtre particulaire.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ALGORITHMS

ACRONYMS

ABC Approximate Bayesian Computation

A2BC Adaptive Approximate Bayesian Computation

AUV Autonomous Underwater Vehicle

EKF Extended Kalman Filter

EnKF Ensemble Kalman Filter

GPS Global Positioning System

i.i.d. independent and identically distributed

IMU Inertial Measurement Unit

IWEnKF Interacting Weighted Ensemble Kalman Filter

KF Kalman Filter

LS Least Square

MAP Maximum A Posteriori

MISE Mean Integrated Square Error

MMSE Minimum Mean Squared Error

PCRB Posterior Cramér-Rao Bound

PF Particle Filter

RBPF Rao-Blackwellised Particle Filter

RMSE Root Mean Square Error

RPF Regularised Particle Filter

SIR-PF Sequential Importance Resampling Particle Filter

TAN Terrain Aided Navigation

UKF Unscented Kalman Filter

WEnKF Weighted Ensemble Kalman Filter

WGS 84 World Geodetic System 1984

NOTATIONS

$\mathbb{N}, \mathbb{N}^\star$	set of natural integers and non-zero natural integers
$\mathbb{R}, \mathbb{R}^+, \mathbb{R}^\star$	set of real numbers, positive real numbers, and non-zero real number
$X \in \mathbb{R}$	scalar value
$\boldsymbol{X} \in \mathbb{R}^d$	vector of $d \in \mathbb{N}^\star$ dimensions
$\boldsymbol{M} \in \mathbb{R}^{n \times m}$	matrix of n lines and m columns
$\|\boldsymbol{X}\|$	Euclidean norm of $\boldsymbol{X}$
$\|\boldsymbol{M}\|_F$	Frobenius norm of $\boldsymbol{M}$
$\boldsymbol{M}^\top$	transpose of the matrix $\boldsymbol{M}$
$\mathbf{1}_d$	identity matrix in $\mathbb{R}^{d \times d}$
$\mathbf{0}_d, \mathbf{0}_{n \times m}$	zero value vector in $\mathbb{R}^d$, zero value matrix in $\mathbb{R}^{n \times m}$
$\mathbb{1}_{\mathcal{S}}$	indicator function on set $\mathcal{S}$
$\boldsymbol{X} \sim p(\boldsymbol{X})$	random vector $\boldsymbol{X}$ is associated to probability density function p
$\mathbb{E}_p(\boldsymbol{X})$	expectation of the random variable $\boldsymbol{X}$ where $\boldsymbol{X} \sim p(\boldsymbol{X})$
$\mathbb{V}\mathrm{ar}_p(\boldsymbol{X})$	variance of the random variable $\boldsymbol{X}$ where $\boldsymbol{X} \sim p(\boldsymbol{X})$
$\mathcal{N}(\boldsymbol{X}; \boldsymbol{M}, \boldsymbol{P})$	Gaussian density with argument $\boldsymbol{X}$, mean $\boldsymbol{M}$ and covariance P
$\wedge$	cross product
$\nabla f(\boldsymbol{X})$	gradient of function f with respect to $\boldsymbol{X}$
$\nabla^2 f(\boldsymbol{X})$	Hessian matrix of function f with respect to $\boldsymbol{X}$
Δt	sampling period
$\boldsymbol{X} \in \mathbb{R}^d$	vector containing the state variables
$b : \mathbb{R}^d \to \mathbb{R}^d$	dynamical function
$\boldsymbol{W}$	process noise
$\boldsymbol{Z}, \boldsymbol{Y} \in \mathcal{R}^{d_y}$	measurements in the continuous and discrete schemes respectively
$h : \mathbb{R}^d \to \mathbb{R}^{d_y}$	observation function
$\boldsymbol{V}$	observation noise
$\mathcal{Y}_t = \{\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_t\}$	stacked vector of all the measurements from time 1 to t

$\mathcal{X}_t = \{\boldsymbol{X}_0, \ldots, \boldsymbol{X}_t\}$	stacked vector of all the state variables from time 0 to t
$p(\boldsymbol{X}_k \vert \mathcal{Y}_k)$	conditional density
$p(\boldsymbol{X}_k \vert \mathcal{Y}_{k-1})$	prior density
$p(\boldsymbol{X}_k \vert \boldsymbol{X}_{k-1})$	transition density
$\boldsymbol{R} \in \mathcal{R}^{d_y \times d_y}$	covariance of the observation noise
$\boldsymbol{Q} \in \mathcal{R}^{d \times d}$	covariance of the process noise
$p(\boldsymbol{Y}_k \vert \boldsymbol{X}_k)$	likelihood
$\boldsymbol{X}^i \in \mathcal{R}^d$	a state sample also called particle
$\{\boldsymbol{X}^i\}_{i=1,\ldots,N}$	N particles
$q(\boldsymbol{X}_k \vert \boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)$	importance density
w^i	normalised weight associated with the particle $\boldsymbol{X}^i$
$\tilde{w}^i$	unnormalised weight associated with the particle $\boldsymbol{X}^i$
δ	Dirac delta function on $\mathbb{R}^d$ equal to 0 on $\mathbb{R}^d \backslash \{\mathbf{0}_d\}$ and to infinity on $\mathbf{0}_d$
N_{th}	resampling threshold
$\mathcal{K}_h : \mathbb{R}^d \to \mathbb{R}$	regularisation kernel
$\mathcal{V}$	asymptotic variance of the unnormalised weights
$\boldsymbol{p} = [p_\lambda, p_\phi, p_h]^\top$	geographical position in the Earth frame where p_λ is the latitude, p_ϕ the longitude and p_h is altitude
$\boldsymbol{V}$	velocity in the Earth frame
$\boldsymbol{\Psi} = [\psi, \theta, \varphi]^\top$	attitude in the Earth frame where ψ is the heading, θ the pitch and φ the roll
$\boldsymbol{Y}^a$	vehicle's acceleration given by the inertial measurement unit in the inertial frame
$\boldsymbol{Y}^\omega$	vehicle's angular rate given by the inertial measurement unit in the inertial frame
$[\boldsymbol{R}_{n2b}]$	attitude matrix
$\boldsymbol{\rho}$	angular velocity of the navigation frame with respect to the Earth frame
$\boldsymbol{\omega}_{ie}$	velocity of the Earth rotation on itself
$\delta\boldsymbol{X}$	inertial measurement unit errors
R_λ	radius of curvature of the Earth in the meridian plane
R_ϕ	large normal of the ellipsoidal Earth
f	vector of proper accelerations projected in the navigation frame
ω_S	vector of Schuler frequency

$\boldsymbol{Y}^b$	multi-beam telemeter measurements
r^i	i-th distance between the vehicle and the seabed
m	number of beams
$\boldsymbol{Y}^g$	atomic gravimeter measurement
$\mathcal{K}_{\varepsilon_k}$	approximate likelihood kernel of bandwidth parameter ε_k
$\boldsymbol{U}$	pseudo-observations
$\mathcal{D} \subset \mathbb{R}^+$	domain of the bandwidth parameter ε_k

1

INTRODUCTION

State estimation aims to retrieve the current state of an evolving system from available measurements. Observation and dynamical models are not perfectly known, which introduces uncertainty in the estimation process. A state estimation algorithm generally outputs a point-wise estimate of the state, associated with a confidence region that is likely to contain the actual state.

Stochastic state estimation algorithms rebuild a conditional state density given the past measurements to determine the most probable current state. In the linear-Gaussian Markovian case, the state density remains Gaussian, and the optimal estimation solution is provided by the Kalman Filter (KF) [44]. However, in the case of non-Gaussian uncertainty, non-linear measurements, and severe ambiguities (i.e., to one measurement occurrence may correspond several possible states), KF and its derivatives fail to estimate the state. Some non-parametric estimation approaches have therefore been introduced, such as the Particle Filter (PF) [34] (a.k.a., sequential Monte Carlo methods), which consists of approximating the state density by a weighted sum of Dirac functions.

More accurate PF derivations were proposed, for example the Regularised Particle Filter (RPF) [61] and the Rao-Blackwellised Particle Filter (RBPF) [17]. RPF consists of smoothing the state density approximation by a mixture of bounded stochastic kernels. RBPF is a variance reduction method for conditionally linear Gaussian models.

Despite these improvements, filters may diverge due to successive Monte Carlo approximations that are introduced for the evaluation of the integrals and at the resampling step. To reduce the variance of the Monte Carlo approximation, two approaches are proposed:

- Guaranteeing the consistency between the likelihood and the prior density (i.e., when the support of the likelihood contains the support of the prior density) is crucial to avoid the degeneracy problem. The proposed Adaptive Approximate Bayesian Computation (A2BC) method ensures the consistency between the likelihood and the prior density by adapting the support of the likelihood. The likelihood is approximated by a kernel density function that spreads out to ensure the consistency without flattening. The information delivered by the measurements is then preserved. The A2BC method reduces the occurrence of the resampling step as the degeneracy problem is mitigated. Thus, the proposed method limits filter divergence.
  The A2BC method is based on Approximate Bayesian Computation (ABC) [16, 24, 55] approaches which address cases where the likelihood is unknown. The likelihood may

be unavailable when the observation model is a rather rough approximation of the true model, or when the likelihood is too expensive to calculate.
The A2BC method can be integrated in any particle-like filter. In this work, it is integrated within RPF and RBPF.

- The choice of the proposal density is crucial as it ensures the correct functioning of particle filter algorithms. Indeed, sequential Monte Carlo methods rely on the generation of particle samples from the proposal density. The proposal density should minimise the Monte Carlo variance. Unfortunately, the optimal proposal density depends on the target density (i.e., the conditional density).
  The Weighted Ensemble Kalman Filter (WEnKF) [69] can be seen as a particle filter with a Gaussian proposal density calculated by a Kalman filter. However, in some cases, the support of the proposal density of the WEnKF does not overlap with the support of the conditional density, which causes a deterioration of the Monte Carlo estimate and may cause the filter to diverge.
  The proposed state estimation algorithm, called Interacting Weighted Ensemble Kalman Filter (IWEnKF), is based on the WEnKF. The IWEnKF introduces a proposal density for which the support is assured to overlap the support of the conditional density. This choice of proposal density minimises the Monte Carlo variance and thus provides a better accuracy and robustness to nonlinearities than the WEnKF does.

The main application is the underwater navigation of an Autonomous Underwater Vehicle (AUV) which consists of retrieving the vehicle's state (position, velocity, and attitude). AUV navigation is often based on Inertial Measurement Unit (IMU) measurements. Although IMU is autonomous and reliable, the measurements it provides are subject to bias and noises that result in a drifting error in the navigation solution. To correct the navigation drift, IMU can be combined with external sensors. In the underwater field, there is a strong demand for autonomy. For example, Global Positioning System (GPS) measurements are unavailable due to the underwater environment, and resurfacing is often excluded for discretion purposes. This has a significant impact on the estimation performance, that must be compensated due to the lack of information by more complex embedded algorithms.

Terrain Aided Navigation (TAN) provides a drift-free navigation tool for AUV operations. This method aims to retrieve the vehicle current state by matching a terrain profile obtained from a sensor with a profile reconstructed from an embedded map of the operation area. This application involves severe measurements ambiguities resulting in the failure of many existing estimation methods. The presence of measurement ambiguities (e.g., to one terrain sensing measurement may correspond several geographical areas of similar terrain profiles) increases the complexity of state estimation, and may prevent estimation algorithms from converging to a unique estimated state with a sufficiently narrow confidence region.

The sensor used in the TAN method is usually a multi-beam telemeter which provides a series of depth measurements along the trajectory [45, 57, 64, 65]. To improve the estimation of the kinematic parameters of the vehicle, another sensor is introduced: the atomic gravimeter. The atomic gravimeter is a promising absolute sensor for underwater navigation. In particular, the quantum gravimetry concept developed by ONERA (cold atom gravimeter [10]) provides

an absolute and accurate gravity measurement. This means that the sensor does not need any calibration and provides a gravity evaluation with a very low level noise. The atomic gravimeter is starting to be used in underwater navigation [62] but has not been combined with other TAN sensors.

The aim of this work is to develop accurate and stable state estimation algorithms based on the TAN method for an AUV equipped with a multi-beam telemeter and an atomic gravimeter.

1.1 THESIS STRUCTURE

The rest of the thesis is organised as follows:

Chapter 2 presents a literature review on the bayesian estimation. The probabilistic framework is introduced for the continuous and discrete schemes. Analytical methods for the linear Gaussian case are recalled.

Chapter 3 presents a literature review on Monte Carlo methods. This chapter presents non-linear non-Gaussian filters such as RPF, RBPF, and WEnKF on which the proposed algorithms are based, and the likelihood-free ABC method. Monte Carlo errors criteria are also introduced.

Chapter 4 introduces the underwater navigation application. Dynamical and observation models under consideration are described, with a focus on the TAN application. Finally, some evaluation criteria are defined.

Chapter 5 introduces a new method called A2BC. Two theoretical results on Monte Carlo errors are proved. A2BC simulations are presented for the underwater TAN application. A2BC filters are tested on a simplified navigation model and on an inertial error model that is in accordance to realistic situations. A2BC filters are compared to their conventional versions namely RPF and RBPF. Simulation results demonstrate that the proposed filters significantly improve the robustness to measurement uncertainties and the accuracy of the state estimates.

Chapter 6 presents a new state estimation algorithm called IWEnKF. Theoretical results leading to an enhanced choice of the proposal density are introduced. IWEnKF simulations are presented for the underwater TAN application. IWEnKF is compared to the filter on which it is based: WEnKF. Simulation results demonstrate that the proposed filter brings more robustness to nonlinearities.

Chapter 7 concludes the thesis and discusses possible directions for future work.

1.2 PUBLICATIONS

The work on which this thesis is based has led to the following publications:

1.2.1 *Journal articles*

[1] Brian DO Anderson, Adrian N Bishop, Pierre Del Moral, and Camille Palmier. "Backward nonlinear smoothing diffusions". In: *Teor. Veroyatnost. i Primenen.* 66.2 (2021), pp. 305–326.

[3] Camille Palmier, Karim Dahia, Nicolas Merlinge, Pierre Del Moral, and Dann Laneuville. "Adaptive Approximate Bayesian Computation Particle Filters - Application to Underwater Terrain Aided Navigation". In: (2021). Submit to: IEEE Transactions on Aerospace and Electronic Systems.

1.2.2 *International Conference Proceedings*

[2] Camille Palmier, Karim Dahia, Nicolas Merlinge, Pierre Del Moral, Dann Laneuville, and Christian Musso. "Adaptive approximate Bayesian computational particle filters for underwater terrain aided navigation". In: *2019 22th International Conference on Information Fusion (FUSION)*. IEEE. 2019, pp. 1–8.

[4] Camille Palmier, Karim Dahia, Nicolas Merlinge, Dann Laneuville, and Pierre Del Moral. "Bathymetry and Atomic Gravimetry Sensor Fusion for Autonomous Underwater Vehicle". In: *2021 23th International Conference on Information Fusion (FUSION)*. IEEE. 2021, pp. 1–8.

[5] Camille Palmier, Karim Dahia, Nicolas Merlinge, Dann Laneuville, and Pierre Del Moral. "Interacting Weighted Ensemble Kalman Filter applied to Underwater Terrain Aided Navigation". In: *2021 American Control Conference (ACC)*. IEEE. 2021, pp. 1541–1546.

2

BAYESIAN ESTIMATION

Let us consider a system (e.g., an autonomous underwater vehicle) whose evolution is described by a set of functions explicitly defined or derived from differential equations. These functions correspond to the evolution model which provides a time description of the system's variables of interest (e.g., position, velocity, attitude etc.). The system does not usually have full knowledge of its state vector. Its decision process is only derived from the information available on the whole vector or some of its components. Since the dynamical model function is most often a simplification of reality, the value of the system state vector obtained as the dynamical model output can rapidly drift far from the actual system state vector. In order to limit this derivation, we can use some measurements provided by sensors which are functions of the state vector components. However, these measurements are obtained via imperfect sensors and thus do not give an exact output. The estimation problem then consists of retrieving the system's state vector description, given an approximate dynamical model and noisy measurements. A specific aspect of the estimation problem addressed here is the search for recursive state estimation. At each time-step, the state vector value is derived from the previous state value, through an update performed by both the dynamical and measurement models.

In this chapter, the estimation problem is first presented in Section 2.1. Then, the probabilistic framework is introduced in Sections 2.2 and 2.3 in the continuous and discrete schemes respectively. Under Gaussian and linear assumptions, this framework leads to Kalman filters and Rauch-Tung-Striebel smoothers (see Section 2.4). Extended observers are presented in Section 2.5 to address cases where the linearity assumption is relaxed.

2.1 PROBLEM STATEMENT

In the continuous time, the signal $\{\boldsymbol{X}_t\}_{t\geq 0} \in \mathbb{R}^d$ contains the variables that define the system (i.e., position, velocity etc.) at a given time as well as its temporal evolution. The state evolution can be defined by the following Itô stochastic differential equation:

$$d\boldsymbol{X}_t = b_t(\boldsymbol{X}_t)\ dt + \sigma_t(\boldsymbol{X}_t)\ d\boldsymbol{W}_t \tag{2.1}$$

where $b_t : \mathbb{R}^d \to \mathbb{R}^d$ is the dynamical function and $d\boldsymbol{W}_t$ a m-dimensional Brownian motion for $m \geq 1$ finite. The process noise $d\boldsymbol{W}_t$ accounts for model uncertainties and unforeseen

disturbances in the state evolution. Let $\boldsymbol{X}_0$ be an initial random variable with absolute moments of any order.

The signal $\boldsymbol{X}_t$ (2.1) is unknown and cannot be directly retrieved due to its initial uncertainty and the signal noise. The state system is partially observed through a series of measurements provided by sensors. In practice, the measurements equation is always expressed in discrete time. We consider a sequence of times $t_0, t_1, \ldots, t_K$. Two consecutive times-steps are separated by a time $\Delta t = t_{k+1} - t_k > 0$. For the sake of simplicity, we set $\boldsymbol{Y}_k \coloneqq \boldsymbol{Y}_{t_k}$ and $\boldsymbol{X}_k \coloneqq \boldsymbol{X}_{t_k}$ where the subscript k takes values in $[0, K]$. Measurements $\{\boldsymbol{Y}_k\}_{k\geq 0} \in \mathbb{R}^{d_y}$ are related to the state by the following observation equation:

$$\boldsymbol{Y}_k = h_k(\boldsymbol{X}_k) + \boldsymbol{V}_k \tag{2.2}$$

where $h_k : \mathbb{R}^d \to \mathbb{R}^{d_y}$ is the observation function and $\boldsymbol{V}_k$ the observation noise, and $\boldsymbol{Y}_0$ is by convention equal to 0. The measurement noise represents the imperfections of the sensors. The observation function h_k may not involve all the state variables. In continuous time, the stochastic differential equation associated with the observation process $\boldsymbol{Z}_t$ is defined by:

$$d\boldsymbol{Z}_t = h_t(\boldsymbol{X}_t)\ dt + d\boldsymbol{V}_t \tag{2.3}$$

where $d\boldsymbol{V}_t$ is a n-dimensional Brownian motion for any $n \geq 1$. The discrete time model (2.2) is recovered by taking $\Delta t \simeq 0$:

$$\boldsymbol{Z}_{t_k} - \boldsymbol{Z}_{t_{k-1}} \simeq_{\Delta t \simeq 0} h_{t_{k-1}}(\boldsymbol{X}_{t_{k-1}})\, \Delta t + \sqrt{\Delta t}\, \boldsymbol{U}_k^{\boldsymbol{Y}} \tag{2.4}$$

for any independent and identically distributed (i.i.d.) Gaussian random variables $\boldsymbol{U}_k^{\boldsymbol{Y}} \sim \mathcal{N}(0, 1)$.

The functions $b_t(\boldsymbol{x})$, $h_t(\boldsymbol{x})$ and $\sigma_t(\boldsymbol{x})$ are assumed to be sufficiently regular so that the filtering problem is well defined. To avoid unnecessary technical discussions, we will assume that the functions $b_u(\boldsymbol{x})$ and $h_u(\boldsymbol{x})$, as well as the diffusion matrix $\sigma_u(\boldsymbol{x})$ are smooth w.r.t. $(u, \boldsymbol{x})$ and that they have uniformly bounded derivatives w.r.t. $\boldsymbol{x}$ of all order on $(u, \boldsymbol{x}) \in [s, t] \times \mathbb{R}^d$, for any $s \leq t$. These conditions ensure that the stochastic differential equations model (equations (2.1) and (2.3)) has a global solution $(\boldsymbol{X}_t, \boldsymbol{Z}_t)$ in the sense of Itô.

The goal of the recursive state estimation is to determine the state $\boldsymbol{X}_t$ at each time-step t given the observations $\mathcal{Z}_T \triangleq (\boldsymbol{Z}_s)_{0\leq s\leq T}$. Depending on the value of T, three problems can be distinguished:

- $T < t$: the prediction problem consists of computing the conditional distribution of the random state $\boldsymbol{X}_t$ given the observations $\mathcal{Z}_T$.
- $T = t$: the filtering problem consists of computing the conditional distribution $\pi_t \coloneqq \pi_{t|t}$ of the random signal states $\boldsymbol{X}_t$ given the observations $\mathcal{Z}_t$.
- $T > t$: the smoothing problem computes the conditional distribution $\pi_{T,t}$ of the random signal states $\boldsymbol{X}_t$ given $\mathcal{Z}_T$.

Knowledge of the conditional density enables us to compute an optimal point-wise state estimate with respect to any criterion. For example, the Minimum Mean Squared Error (MMSE) estimate is the conditional mean of $\boldsymbol{X}_t$ [7]:

$$\hat{\boldsymbol{X}}_{t|t}^{MMSE} \coloneqq \mathbb{E}_p(\boldsymbol{X}_t|\mathcal{Z}_t) = \int \boldsymbol{X}_t \; p(\boldsymbol{X}_t|\mathcal{Z}_t) d\boldsymbol{X}_t \tag{2.5}$$

Whereas the Maximum A Posteriori (MAP) is the maximum of the conditional density $p(\boldsymbol{X}_t|\mathcal{Z}_t)$ [31]:

$$\hat{\boldsymbol{X}}_{t|t}^{MAP} \coloneqq \arg\max_{\boldsymbol{X}_k} p(\boldsymbol{X}_t|\mathcal{Z}_t) \tag{2.6}$$

If the conditional density is unimodal, i.e., it has a unique global maximum, the MAP is an optimal estimator. We could also cite the Least Square (LS) [50] estimator which consists of maximising a Gaussian likelihood $p(\boldsymbol{Z}_t|\boldsymbol{X}_t) \propto \exp\left(\frac{1}{2}(\boldsymbol{Z}_t - h_t(\boldsymbol{X}_t))^\top \boldsymbol{R}_t^{-1}(\boldsymbol{Z}_t - h_t(\boldsymbol{X}_t))\right)$:

$$\hat{\boldsymbol{X}}_{t|t}^{LS} \coloneqq \arg\min_{\boldsymbol{X}_t} \Big((\boldsymbol{Z}_t - h_t(\boldsymbol{X}_t))^\top \boldsymbol{R}_t^{-1}(\boldsymbol{Z}_t - h_t(\boldsymbol{X}_t))\Big) \tag{2.7}$$

where $\boldsymbol{R}_t$ is the covariance of the observation noise. A confidence region of the state estimation (e.g., the variance) can also be obtained. In general, the recursive propagation of the conditional distribution cannot be determined analytically. The implementation of the solution requires the storage of the entire density which is, in general, an infinite dimensional vector. Thus, instead of approaching the optimal filter density at any point in space, only a few moments (typically the mean and variance) as well as the location of the modes (when the density is multimodal) are really required.

An iterative estimation algorithm often consists of two steps: the prediction step, which accounts for the dynamical transition, and the correction step, which accounts for the measurements if available.

2.2 CONTINUOUS FILTERING AND SMOOTHING

The probabilistic approach is described, with emphasis on optimality in estimation. The equations of evolution for the conditional probability density function are developed in both continuous and discrete times for the filtering and smoothing problems. A detailed derivation of the probabilistic approach can be found in [25], [78] and [6] and references therein.

2.2.1 *Optimal Filter*

The probabilistic framework aims to estimate the most probable current state $\boldsymbol{X}_t$ given all the past measurements $\mathcal{Z}_t$ by estimating the conditional density π_t.

The conditional density can be computed in an iterative way via the Kolmogorov forward equation (2.8) [53] and Bayes' rule (2.10). The equations consist of a theoretical filter named optimal filter. It represents the most accurate way of describing the conditional density function evolution, given a trajectory of measurements and their associated uncertainties.

The evolution of the optimal filter is composed of two steps:

- The prediction step determines the evolution of the prior density via the Kolmogorov forward equation:

$$\pi_t = \mathcal{L}\text{aw}(\boldsymbol{X}_t) \Leftrightarrow d\pi_t(f) = \pi_t(L_t(f))\, dt \tag{2.8}$$

  where $\pi(f) \coloneqq \int \pi(d\boldsymbol{x})\, f(\boldsymbol{x})$ and f stands for any twice differentiable test function with bounded derivatives. The generator L_t of the signal $\boldsymbol{X}_t$ is given by the second order differential operator:

$$L_t(f)(\boldsymbol{x}) \coloneqq \nabla f(\boldsymbol{x})^\top b_t(\boldsymbol{x}) + \frac{1}{2}\, \text{Tr}(\nabla^2 f(\boldsymbol{x}) \alpha_t(\boldsymbol{x})) \tag{2.9}$$

  where $\alpha_t(\boldsymbol{x}) = \sigma_t(\boldsymbol{x})\sigma_t(\boldsymbol{x})^\top$, Tr is the trace operator. ∇f is the column gradient and $\nabla^2 f$ the Hessian matrix of the function f.

- Roughly speaking, when a new measurement $d\boldsymbol{Z}_t$ is available, it is associated with a sensor noise distribution conditional to the current state, which is called the likelihood. The likelihood quantifies the distribution of the measurement according to the actual state. The correction step determines the conditional density of the state with respect to the prior density (2.8) and the likelihood $p(d\boldsymbol{Z}_t|\boldsymbol{X}_t)$. From Bayes' law, this discrete time reasoning yields:

$$p(\boldsymbol{X}_t|\mathcal{Z}_t, d\boldsymbol{Z}_t) = \frac{p(d\boldsymbol{Z}_t|\boldsymbol{X}_t)\; p(\boldsymbol{X}_t|\mathcal{Z}_t)}{p(d\boldsymbol{Z}_t|\mathcal{Z}_t)} \tag{2.10}$$

  The denominator can be evaluated using the Markovian property of the model:

$$p(d\boldsymbol{Z}_t|\mathcal{Z}_t) = \int p(d\boldsymbol{Z}_t|\boldsymbol{X}_t)\; p(\boldsymbol{X}_t|\mathcal{Z}_t)\; d\boldsymbol{X}_t \tag{2.11}$$

The recurrence relations (2.8)-(2.10) form the basis for the optimal Bayesian solution. The combination of the prediction and correction steps leads to the following generator of the conditional distribution:

$$d\pi_t(f) = \pi_t(L_t(f))\, dt + \pi_t(f(h_t - \pi_t(h_t)))^\top \boldsymbol{R}_t^{-1}(d\boldsymbol{Z}_t - \pi_t(h_t)) \tag{2.12}$$

where $\boldsymbol{R}_t\, dt \coloneqq \mathbb{E}(d\boldsymbol{V}_t\, d\boldsymbol{V}_t^\top)$ is the covariance of the observation noise. As in equation (2.8), the function f stands for any twice differentiable function with bounded derivatives. Equation (2.12) is called the Kushner-Stratonovich equation.

The optimal filter is illustrated in Figure 2.1.

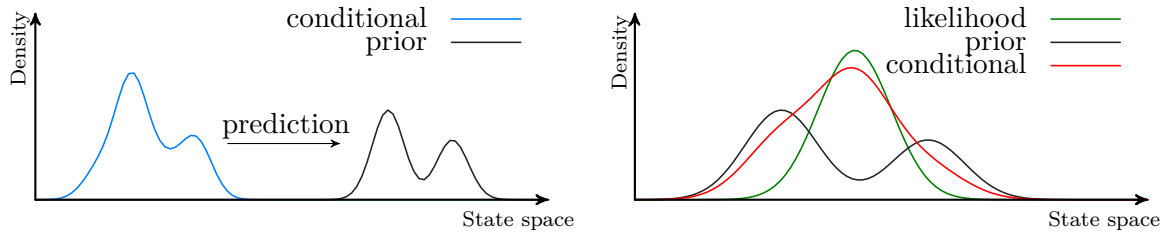

(a) Prediction step: the previous conditional density is convoluted with the transition density. This results in the prior density, for which the support is usually larger than the previous conditional density.

(b) Correction step: when the measurement is available, the prior density and the likelihood lead to the determination of the conditional density.

Figure 2.1: Probabilistic iterative prediction-correction scheme.

2.2.2 *Optimal Smoother*

The smoothing aims to compute the conditional distribution $\pi_{t,s}$ of the random signal states $\boldsymbol{X}_s$ given $\mathcal{Z}_t$ with $t \geq s$. We assume for any $t > 0$ that the conditional distribution π_t has a positive density $p_t \coloneqq d\pi_t/d\lambda$ w.r.t. the Lebesgue measure on $\mathbb{R}^d$. In addition, $p_u(\boldsymbol{x})$ and its derivative $\nabla p_u(\boldsymbol{x})$ are uniformly bounded w.r.t. $(\boldsymbol{u}, \boldsymbol{x}) \in [s,t] \times \mathbb{R}^d$, for any given $s > 0$, almost surely w.r.t. the distribution of the observation process.

From the backward Itô integration formula, we obtain:

$$\mathcal{X}_{t,s}(\boldsymbol{x}) = \boldsymbol{x} + \int_s^t \left(p_u(\mathcal{X}_{t,u}(\boldsymbol{x}))^{-1} \mathrm{div}_{\alpha_u(p_u)}(\mathcal{X}_{t,u}(\boldsymbol{x})) - b_u(\mathcal{X}_{t,u}(\boldsymbol{x})) \right) du + \int_s^t \sigma_u(\mathcal{X}_{t,u}(\boldsymbol{x})) d\mathcal{W}_u \tag{2.13}$$

with the terminal condition $\mathcal{X}_{t,t}(\boldsymbol{x}) = \boldsymbol{x}$, and where $\mathcal{W}_t \in \mathbb{R}^m$ denotes a m-dimensional Brownian motion independent of the observation. With a sufficiently smooth function $f : \mathbb{R}^d \to \mathbb{R}$, we note $\mathrm{div}_{\alpha_t}(f)$ the α_t-divergent d-column vector operator with j-th entry given by the formula:

$$\mathrm{div}_{\alpha_t}(f)(\boldsymbol{x})^j \coloneqq \sum_{1 \leq i \leq d} \partial_{\boldsymbol{x}_i} \left(\alpha_t^{i,j}(\boldsymbol{x}) f(\boldsymbol{x}) \right) \tag{2.14}$$

Let $\mathcal{X}_t$ be a random variable with the distribution π_t, for $t > 0$. The backward process $\mathcal{X}_{t,s}(\mathcal{X}_t)$ is distributed according to the smoothing distribution $\pi_{t,s}$ for any $s \geq t$ whenever the terminal condition $\mathcal{X}_{t,t}(\mathcal{X}_t) = \mathcal{X}_t$ is distributed according to the filtering distribution π_t.

The generator $\mathcal{L}_{s,\pi_s}$ of the signal $\mathcal{X}_{t,s}(\mathcal{X}_t)$ is given by the second order differential operator:

$$\mathcal{L}_{s,\pi_s}(f) = \sum_{1 \leq j \leq d} \left(-b_s^j + \frac{1}{p_s} \mathrm{div}_{\alpha_s}(p_s)^j \right) \partial_{\boldsymbol{x}_j} f + \frac{1}{2} \sum_{1 \leq i,j \leq d} \alpha_s^{i,j} \partial_{\boldsymbol{x}_i \boldsymbol{x}_j} f \tag{2.15}$$

The smoothing generator (2.15) is similar to the filtering generator (2.9) except for the sign of the drift b_t and the presence of the forward conditional density p_s.

The backward evolution equation is given by:

$$\pi_{t,s} = \mathcal{L}\text{aw}(\mathcal{X}_{t,s}) \Leftrightarrow \partial_s \pi_{t,s}(f) = -\pi_{t,s}(\mathcal{L}_{s,\pi_s}(f)) \tag{2.16}$$

with the terminal condition $\pi_{t,t} = \pi_t$.

2.3 DISCRETE FILTERING AND SMOOTHING

Similarly, the optimal filter and smoother equations can be written in discrete time. We consider the state system $\boldsymbol{X}_t$ at a discrete sequence of times $t_0, t_1, \ldots, t_K$. We set $\boldsymbol{X}_k := \boldsymbol{X}_{t_k}$ where the subscript k takes its values in $[0, K]$. Two consecutive time-steps are separated by a time $\Delta t = t_{k+1} - t_k > 0$. The state evolution can be defined by the following equation:

$$\boldsymbol{X}_k = b_k(\boldsymbol{X}_{k-1}) + \boldsymbol{W}_k \tag{2.17}$$

where $\boldsymbol{W}_k$ is the signal noise. The discrete time model (2.17) comes from the continuous model (2.1) by taking $\Delta t \simeq 0$:

$$\boldsymbol{X}_{t_k} - \boldsymbol{X}_{t_{k-1}} \simeq_{\Delta t \simeq 0} b_{t_{k-1}}(\boldsymbol{X}_{t_{k-1}})\,\Delta t + \sqrt{\Delta t}\,\sigma_{t_{k-1}}(\boldsymbol{X}_{t_{k-1}})\,\boldsymbol{U}_k^{\boldsymbol{X}} \tag{2.18}$$

for any i.i.d. Gaussian random variables $\boldsymbol{U}_k^{\boldsymbol{X}} \sim \mathcal{N}(0,1)$.

The discrete measurement equation is given by (2.2) and is repeated here:

$$\boldsymbol{Y}_k = h_k(\boldsymbol{X}_k) + \boldsymbol{V}_k \tag{2.19}$$

The initial state $\boldsymbol{X}_0$ of known density $p(\boldsymbol{X}_0)$ is independent of the two noises $\boldsymbol{W}_k$ and $\boldsymbol{V}_k$. The noise sequences $\boldsymbol{W}_k$ and $\boldsymbol{V}_k$ are both assumed white and mutually independent.

2.3.1 *Optimal filter*

The optimal filter aims to estimate the conditional density $p(\boldsymbol{X}_k|\mathcal{Y}_k)$ of the random signal states $\boldsymbol{X}_k$ given the sigma-field $\mathcal{Y}_k = \sigma(\boldsymbol{Y}_s, s \leq t)$ generated by the observations. The evolution of the optimal filter is composed of two steps:

- The prediction step determines the prior density $p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})$ with respect to the transition density $p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1})$ and the previous conditional density $p(\boldsymbol{X}_{k-1}|\mathcal{Y}_{k-1})$ via the Chapman-Kolmogorov equation:

$$p(\boldsymbol{X}_k|\mathcal{Y}_{k-1}) = \int p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1})\; p(\boldsymbol{X}_{k-1}|\mathcal{Y}_{k-1})\; d\boldsymbol{X}_{k-1} \tag{2.20}$$

  This equation is based on the assumptions on the noise sequences and on Bayes' rule, which allows us to prove that all measurements are independent given the states and

that the state evolution process is Markovian, i.e., the state $\boldsymbol{X}_k$ depends only on the previous state $\boldsymbol{X}_{k-1}$. By denoting $\mathcal{X}_k = \{\boldsymbol{X}_1, \dots, \boldsymbol{X}_k\}$ the stacked vector of all the states up to time t, $\boldsymbol{X}_k$ is a Markov process of order one if:

$$p(\mathcal{X}_k) = \prod_{t=0}^{k} p(\boldsymbol{X}_t|\boldsymbol{X}_{t-1}) \tag{2.21}$$

where $p(\boldsymbol{X}_0|\boldsymbol{X}_{-1}) = p(\boldsymbol{X}_0)$. The conditional density of the measurements given the states is independent if:

$$p(\mathcal{Y}_k|\mathcal{X}_k) = \prod_{t=0}^{k} p(\boldsymbol{Y}_t|\boldsymbol{X}_t) \tag{2.22}$$

- When a new measurement $\boldsymbol{Y}_k$ is available, it is associated with a sensor noise density $p(\boldsymbol{Y}_k|\boldsymbol{X}_k)$ conditionally to the current state. Using the assumptions (2.21)-(2.22) we derive the correction step. From Bayes' law, we obtain:

$$p(\boldsymbol{X}_k|\mathcal{Y}_k) = \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\ p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{\int p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\ p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\ d\boldsymbol{X}_k} \tag{2.23}$$

2.3.2 *Optimal Smoother*

The smoothing problem consists in estimating the conditional density $p(\boldsymbol{X}_k|\mathcal{Y}_K)$ where $k \in [0, K]$. The backward recursive equations are given by:

$$p(\boldsymbol{X}_{k+1}|\mathcal{Y}_k) = \int p(\boldsymbol{X}_{k+1}|\boldsymbol{X}_k)\, p(\boldsymbol{X}_k|\mathcal{Y}_k)\, d\boldsymbol{X}_k \tag{2.24}$$

$$p(\boldsymbol{X}_k|\mathcal{Y}_K) = p(\boldsymbol{X}_k|\mathcal{Y}_k) \int \frac{p(\boldsymbol{X}_{k+1}|\boldsymbol{X}_k)\, p(\boldsymbol{X}_{k+1}|\mathcal{Y}_K)}{p(\boldsymbol{X}_{k+1}|\mathcal{Y}_k)}\, d\boldsymbol{X}_{k+1} \tag{2.25}$$

where $p(\boldsymbol{X}_k|\mathcal{Y}_k)$ is the filtering density. The recursion is started from the final step, when the filtering and smoothing densities are the same: $p(\boldsymbol{X}_k|\mathcal{Y}_K)$.

The next section discusses the linear Gaussian case leading to the Kalman Filter (KF) and Rauch-Tung-Striebel smoother. The linear assumption is then relaxed to introduce the Extended Kalman Filter (EKF) in Section 2.5.

2.4 THE LINEAR AND GAUSSIAN CASE

In the linear Gaussian case, the conditional distribution can be exactly and completely characterised by its mean and covariance. This section restates the KF and the Gaussian smoother equations [44].

2.4.1 *Continuous Kalman-Bucy Filter and Rauch-Tung-Striebel Smoother*

The state evolution and measurements equation are assumed to be linear Gaussian models defined by:

$$\begin{cases} d\boldsymbol{X}_t = \boldsymbol{B}_t\, \boldsymbol{X}_t\, dt + \boldsymbol{\Sigma}_t\, d\boldsymbol{W}_t \\ d\boldsymbol{Z}_t = \boldsymbol{H}_t\, \boldsymbol{X}_t\, dt + d\boldsymbol{V}_t \end{cases} \tag{2.26}$$

where $\boldsymbol{B}_t \in \mathbb{R}^{d\times d}$ and $\boldsymbol{H}_t \in \mathbb{R}^{d_y\times d}$ are known matrices. $\boldsymbol{W}_t$ and $\boldsymbol{V}_t$ are Brownian motions of covariances $\mathbb{E}(d\boldsymbol{W}_t\, d\boldsymbol{W}_t^\top) = \boldsymbol{Q}_t\, dt > 0$ and $\mathbb{E}(d\boldsymbol{V}_t\, d\boldsymbol{V}_t^\top) = \boldsymbol{R}_t\, dt > 0$ respectively. The initial random variable $\boldsymbol{X}_0$ is assumed to be Gaussian of mean $\hat{\boldsymbol{X}}_0$ and covariance matrix P_0.

2.4.1.1 *Continuous Kalman-Bucy Filter*

The conditional density π_t is Gaussian and can be fully described by its expectancy $\hat{\boldsymbol{X}}_t = \mathbb{E}[\boldsymbol{X}_t|\mathcal{Z}_t]$ and its covariance $\boldsymbol{P}_t = \mathbb{E}[(\boldsymbol{X}_t - \hat{\boldsymbol{X}}_t)(\boldsymbol{X}_t - \hat{\boldsymbol{X}}_t)^\top]$. The expectancy $\hat{\boldsymbol{X}}_t$ is equal to the theoretical MAP estimator (2.6). The Kalman estimator is optimal in the sense that $\hat{\boldsymbol{X}}_t$ and $\boldsymbol{P}_t$ are equal to the conditional mean and the covariance of the optimal filter Gaussian conditional density.

The optimal filter π_t is a Gaussian distribution with mean and covariance satisfying the Kalman-Bucy and the Ricatti equations:

$$\begin{cases} d\hat{\boldsymbol{X}}_t = \boldsymbol{B}_t\, \hat{\boldsymbol{X}}_t\, dt + \boldsymbol{P}_t\boldsymbol{H}_t^\top \boldsymbol{R}_t^{-1} \left(d\boldsymbol{Z}_t - \boldsymbol{H}_t\hat{\boldsymbol{X}}_t\, dt\right) \\ \partial_t \boldsymbol{P}_t = \boldsymbol{B}_t\boldsymbol{P}_t + \boldsymbol{P}_t\boldsymbol{B}_t^\top + \boldsymbol{\Sigma}_t\boldsymbol{Q}_t{\boldsymbol{\Sigma}_t}^\top - \boldsymbol{P}_t\boldsymbol{H}_t^\top\boldsymbol{R}_t^{-1}\boldsymbol{H}_t\boldsymbol{P}_t \end{cases} \tag{2.27}$$

where ∂_t denotes $\frac{d}{dt}$.

The matrix $\boldsymbol{P}_t\boldsymbol{H}_t^\top\boldsymbol{R}_t^{-1}$ is called the Kalman gain. The Kalman gain is multiplied by the innovation $\left(d\boldsymbol{Z}_t - \boldsymbol{H}_t\hat{\boldsymbol{X}}_t\, dt\right)$ (i.e., the difference between the observed measurements and the predicted measurements). The correction step can be interpreted as follows: a high confidence in the previous estimates ($\|\boldsymbol{P}_t\|$ "small") and a doubt in the measurements ($\|\boldsymbol{R}_t\|$ "large") imply a small Kalman gain and a weakly corrected estimate. Alternatively, a doubt in the previous estimates and a high confidence in the measurements lead to a large Kalman gain and a strongly corrected estimate. The measurements thus have an important role in the final value of the estimate.

2.4.1.2 *Continuous Rauch-Tung-Striebel Smoother*

For any $t \geq s$, $\pi_{t,s}$ is the distribution of $\mathcal{X}_{t,s}(\mathcal{X}_t)$. Since the process is linear, the distribution $\pi_{t,s}$ is Gaussian of mean $\hat{\boldsymbol{X}}_{t,s}$ and covariance $\boldsymbol{P}_{t,s}$.

We deduce the Rauch-Tung-Striebel smoothing equations [75]. For any $t \geq s$, the parameters $(\hat{\boldsymbol{X}}_{t,s}, \boldsymbol{P}_{t,s})$ satisfy the backward evolution equations:

$$\begin{cases} \partial_s \hat{\boldsymbol{X}}_{t,s} = \boldsymbol{B}_s \hat{\boldsymbol{X}}_{t,s} + \boldsymbol{\Sigma}_s \boldsymbol{Q}_s \boldsymbol{\Sigma}_s{}^\top \boldsymbol{P}_s^{-1} \left(\hat{\boldsymbol{X}}_{t,s} - \hat{\boldsymbol{X}}_s \right) \\ \partial_s \boldsymbol{P}_{t,s} = \left(\boldsymbol{B}_s + \boldsymbol{\Sigma}_s \boldsymbol{Q}_s \boldsymbol{\Sigma}_s{}^\top \boldsymbol{P}_s^{-1} \right) \boldsymbol{P}_{t,s} + \boldsymbol{P}_{t,s} \left(\boldsymbol{B}_s + \boldsymbol{\Sigma}_s \boldsymbol{Q}_s \boldsymbol{\Sigma}_s{}^\top \boldsymbol{P}_s^{-1} \right)^\top - \boldsymbol{\Sigma}_s \boldsymbol{Q}_s \boldsymbol{\Sigma}_s{}^\top \end{cases} \tag{2.28}$$

where $\hat{\boldsymbol{X}}_s$ and $\boldsymbol{P}_s$ are given by the Kalman-Bucy filter and with the terminal condition $(\hat{\boldsymbol{X}}_{t,t}, \boldsymbol{P}_{t,t}) = (\hat{\boldsymbol{X}}_t, \boldsymbol{P}_t)$.

A novel interpretation of the smoothing solution in terms of a nonlinear diffusion stochastic flow (in the spirit of McKean-Vlasov-type processes) is proposed in Anderson et al. [2]. This approach combines forward and backward Itô formulas for stochastic transport semigroups with a recent backward version of the Itô-Ventzell formula. The article is provided in Appendix a.

2.4.2 *Discrete Kalman Filter and Rauch-Tung-Striebel Smoother*

The state evolution and measurements equation are assumed to be linear Gaussian models defined by:

$$\begin{cases} \boldsymbol{X}_k = \boldsymbol{B}_{k-1} \, \boldsymbol{X}_{k-1} + \boldsymbol{\Sigma}_k \boldsymbol{W}_k \\ \boldsymbol{Y}_k = \boldsymbol{H}_k \, \boldsymbol{X}_k + \boldsymbol{V}_k \end{cases} \tag{2.29}$$

where $\boldsymbol{B}_{k-1} \in \mathbb{R}^{d \times d}$ and $\boldsymbol{H}_k \in \mathbb{R}^{d_y \times d}$ are known matrices. The process and observation noises are assumed to be mutually independent zero-mean white Gaussian noises of covariances $\boldsymbol{Q}_k \in \mathbb{R}^{d \times d}$ and $\boldsymbol{R}_k \in \mathbb{R}^{d_y \times d_y}$ respectively. The initial law of the state is Gaussian:

$$p(\boldsymbol{X}_0) \sim \mathcal{N}(\boldsymbol{X}_0; \hat{\boldsymbol{X}}_0, \boldsymbol{P}_0) \tag{2.30}$$

where the mean $\hat{\boldsymbol{X}}_0 \in \mathbb{R}^d$ is the initial guess for the state value, and $\boldsymbol{P}_0 \in \mathbb{R}^{d \times d}$ the initial state covariance. The notation $\mathcal{N}(\boldsymbol{X}; \boldsymbol{M}, \boldsymbol{P})$ represents a Gaussian density with argument $\boldsymbol{X}$, mean $\boldsymbol{M}$, and covariance $\boldsymbol{P}$; such that:

$$\mathcal{N}(\boldsymbol{X}; \boldsymbol{M}, \boldsymbol{P}) \triangleq \frac{1}{|2\pi \boldsymbol{P}|^{1/2}} exp \left\{ -\frac{1}{2} (\boldsymbol{X} - \boldsymbol{M})^\top \boldsymbol{P}^{-1} (\boldsymbol{X} - \boldsymbol{M}) \right\} \tag{2.31}$$

where $\boldsymbol{P}^\top$ refers to the transpose of a matrix P.

2.4.2.1 *Discrete Kalman Filter*

It is well known [40] that the conditional density is Gaussian $p(\boldsymbol{X}_k | \mathcal{Y}_k) \sim \mathcal{N}(\boldsymbol{X}_k; \hat{\boldsymbol{X}}_k, \boldsymbol{P}_k)$ and can be fully described by its expectancy $\hat{\boldsymbol{X}}_k$ and its covariance $\boldsymbol{P}_k$.

KF's equations are presented in Algorithm 2.1.

Algorithm 2.1 Kalman Filter

Initialisation: The initial estimate is $\hat{\boldsymbol{X}}_0$ and the initial covariance is $\boldsymbol{P}_0$.
for each time-step k **do**
  Prediction:
  $\hat{\boldsymbol{X}}_{k|k-1} = \boldsymbol{B}_{k-1}\ \hat{\boldsymbol{X}}_{k-1}$
  $\boldsymbol{P}_{k|k-1} = \boldsymbol{B}_k \boldsymbol{P}_{k-1} \boldsymbol{B}_k^\top + \boldsymbol{\Sigma}_k \boldsymbol{Q}_k \boldsymbol{\Sigma}_k^\top$
  Correction:
  $\hat{\boldsymbol{X}}_k = \hat{\boldsymbol{X}}_{k|k-1} + \boldsymbol{K}_k(\boldsymbol{Y}_k - \boldsymbol{H}_k\ \hat{\boldsymbol{X}}_{k|k-1})$
  $\boldsymbol{P}_k = (\boldsymbol{I}_d - \boldsymbol{K}_k \boldsymbol{H}_k)\boldsymbol{P}_{k|k-1}$
  where $\boldsymbol{K}_k$ is the Kalman gain:
  $\boldsymbol{K}_k = \boldsymbol{P}_{k|k-1}\boldsymbol{H}_k^\top \left(\boldsymbol{H}_k \boldsymbol{P}_{k|k-1} \boldsymbol{H}_k^\top + \boldsymbol{R}_k\right)^{-1}$
end for
Return $\hat{\boldsymbol{X}}_k$ and $\boldsymbol{P}_k$, $\forall k$.

2.4.2.2 *Discrete Rauch-Tung-Striebel Smoother*

Similarly, under the linear Gaussian assumptions, the conditional density $p(\boldsymbol{X}_k|\mathcal{Y}_K)$ is a Gaussian of mean $\hat{\boldsymbol{X}}_{k|K}$ and covariance $\boldsymbol{P}_{k|K}$ given by the following equations:

$$\hat{\boldsymbol{X}}_{k|K} = \hat{\boldsymbol{X}}_k + \boldsymbol{S}_k \left(\hat{\boldsymbol{X}}_{k+1|K} - \boldsymbol{B}_{k+1}\hat{\boldsymbol{X}}_k\right) \tag{2.32}$$

$$\boldsymbol{S}_k = \boldsymbol{P}_k \boldsymbol{B}_{k+1}^\top \left(\boldsymbol{B}_{k+1}\boldsymbol{P}_k\boldsymbol{B}_{k+1}^\top + \boldsymbol{\Sigma}_{k+1}\boldsymbol{Q}_{k+1}\boldsymbol{\Sigma}_{k+1}^\top\right)^{-1} = \boldsymbol{P}_k \boldsymbol{B}_{k+1}^\top \boldsymbol{P}_{k+1|k}^{\ -1} \tag{2.33}$$

$$\boldsymbol{P}_{k|K} = \boldsymbol{P}_k + \boldsymbol{S}_k \left(\boldsymbol{P}_{k+1|K} - \boldsymbol{P}_{k+1|k}\right) \boldsymbol{S}_k^\top \tag{2.34}$$

where $\hat{\boldsymbol{X}}_k$, $\boldsymbol{P}_k$ and $\boldsymbol{P}_{k+1|k}$ are given by the Kalman filter and with the terminal condition $(\hat{\boldsymbol{X}}_{K|K}, \boldsymbol{P}_{K|K}) = (\hat{\boldsymbol{X}}_K, \boldsymbol{P}_{K|K})$.

2.5 EXTENDED OBSERVER ALGORITHMS

EKF [1] is an extension of KF to non-linear Gaussian dynamical models. The conditional density is approximated by a Gaussian whose moments are obtained via linearisation.

In continuous time, the equations of the model are (2.1) and (2.3). Equations (2.17) and (2.2) described the model for the discrete time case.

EKF locally linearises the non-linear functions $b_k : \mathbb{R}^d \to \mathbb{R}^d$ and $h_k : \mathbb{R}^{d_y} \to \mathbb{R}^d$ at each time-step. The functions are approximated by the first term in their Taylor series expansion. The Jacobians ∇b_k and ∇h_k have to be worked out analytically. Thus, EKF cannot be applied if the non-linear functions b_k and h_k are discontinuous.

The equations of the continuous time EKF are:

$$\begin{cases} d\hat{\boldsymbol{X}}_t = b_t(\hat{\boldsymbol{X}}_t)dt + \boldsymbol{P}_t\, \nabla h_t(\hat{\boldsymbol{X}}_t)^\top\, \boldsymbol{R}_t^{-1} \left(d\boldsymbol{Z}_t - h_t(\hat{\boldsymbol{X}}_t)dt\right) \\ \partial_t \boldsymbol{P}_t = \nabla b_t(\hat{\boldsymbol{X}}_t)\, \boldsymbol{P}_t + \boldsymbol{P}_t\, \nabla b_t(\hat{\boldsymbol{X}}_t)^\top + \boldsymbol{\Sigma}_t \boldsymbol{Q}_t \boldsymbol{\Sigma}_t{}^\top - \boldsymbol{P}_t\, \nabla h_t(\hat{\boldsymbol{X}}_t)^\top\, \boldsymbol{R}_t^{-1} \nabla h_t(\hat{\boldsymbol{X}}_t)\, \boldsymbol{P}_t \end{cases} \tag{2.35}$$

In the discrete time case, the system is linearised:

$$\begin{cases} \boldsymbol{X}_k = b_k(\hat{\boldsymbol{X}}_{k-1}) + \nabla b_k(\boldsymbol{X}_{k-1} - \hat{\boldsymbol{X}}_{k-1}) + \boldsymbol{W}_k \\ \boldsymbol{Y}_k = h_k(\hat{\boldsymbol{X}}_{k|k-1}) + \nabla h_k(\boldsymbol{X}_{k-1} - \hat{\boldsymbol{X}}_{k|k-1}) + \boldsymbol{V}_k \end{cases} \tag{2.36}$$

The use of KF for this system provides the equations of EKF at first order.

Similar linearisations are performed to extend the Rauch-Tung-Striebel smoother.

EKF assumes that local linearisation of the equations is a sufficient description of non-linearity. However, if the non-linearity is severe, the non-Gaussianity of the true conditional density will be pronounced as it can be multimodal. In such cases, the performance of EKF will be degraded significantly. Equivalent filters can be deduced from linearisation at higher orders, but they are not used in practice.

The Unscented Kalman Filter (UKF) [87] has been proposed to make EKF more robust. UKF consists of describing the Gaussian state density approximation by a set of deterministically chosen sigma-points in the state space. Each sigma-points is independently propagated through the non-linear state dynamic, and corrected through the non-linear observation equation. The sigma-points are chosen such that their mean, covariance and possibly also higher order moments match the Gaussian state density. The Gaussian state density is computed by a weighted sum of the sigma-points. It results in an algorithm more robust to non-linearity.

2.6 SUMMARY

Estimation schemes were presented in the above sections, leading to a selection of several emblematic approaches. To summarise, probabilistic estimation (Sections 2.2 and 2.3) makes it possible to iteratively estimate a conditional state density knowing a trajectory of measurements. The Kalman filter (Section 2.4) is optimal for the Gaussian linear case and has minimum variance for non-Gaussian linear cases. It was extended to non-linear cases (EKF, UKF in Section 2.5) but these approaches are sub-optimal and suffer from their lack of guarantees. To cope with non-linear cases, Monte Carlo approaches were introduced which empirically describe the non-linear propagation and update of the conditional density. A literature review on Monte Carlo methods is provided in Chapter 3.

3

MONTE CARLO METHODS

Although efficient under specific assumptions, the Kalman Filter (KF) and its extension (e.g., Extended Kalman Filter (EKF), Unscented Kalman Filter (UKF)) suffer from important disadvantages. These methods are based on strong assumptions about the densities involved which are rarely encountered in practice. To tackle a larger class of problems, sequential Monte Carlo methods (a.k.a., Particle Filters (PFs)) were proposed. PF consists of stochastically approximating the state density with a representative sample of possible states. It is able to tackle model non-linearities and is not limited to Gaussian densities.

After a reminder on Monte Carlo estimation (see Section 3.1), this chapter presents a way of deriving the optimal filter in non-linear and non-Gaussian cases. A filter based on Monte Carlo sampling is described in Section 3.2: the Ensemble Kalman Filter (EnKF) [28]. PF scheme [34] is first introduced in Section 3.3.1. PF approximates the state density with a mixture of weighted Dirac deltas. Regularised Particle Filter (RPF) [61], Rao-Blackwellised Particle Filter (RBPF) [17], and Weighted Ensemble Kalman Filter (WEnKF) [69] are then introduced. RPF is based on kernel estimation approaches [81] which brings more accuracy by considering mixtures of weighted bounded kernels. RBPF is a variance reduction method for conditionally linear Gaussian models. WEnKF is based on EnKF and consists in obtaining an exact formulation of the conditional density by adding a weighting step based on the importance sampling PF principle. The uncertainties involve in the probabilistic measurements density function may be unknown. This hypothesis yields the approximate Bayesian computation framework (Section 3.4). Section 3.5 quantifies the Monte Carlo errors. Finally, Section 3.6 gives an overview on Monte Carlo methods.

3.1 MONTE CARLO ESTIMATION

Let $\boldsymbol{X}$ be a random variable on $\mathbb{R}^d$ distributed according to some probability measure $p(\boldsymbol{X})\,d\boldsymbol{X}$ and $(\boldsymbol{X}^i)_{i=1,\dots,N}$ a set of independent random variables on $\mathbb{R}^d$ with the same distribution as $\boldsymbol{X}$. For any bounded function $\mu : \mathbb{R}^d \to \mathbb{R}$, the mean of $\mu(\boldsymbol{X})$ is given by:

$$\mathbb{E}_p\left(\mu(\boldsymbol{X})\right) = \int \mu(\boldsymbol{X})\, p(\boldsymbol{X}) d\boldsymbol{X} \tag{3.1}$$

Monte Carlo methods approximate the expectancy with a set of independent and identically distributed (i.i.d.) random variables $(\boldsymbol{X}^i)_{i=1,\dots,N}$. The Monte Carlo estimator is defined by:

$$\bar{\mu}(\boldsymbol{X}) \coloneqq \frac{1}{N}\sum_{i=1}^{N}\mu(\boldsymbol{X}^i), \quad \boldsymbol{X}^i \sim p \tag{3.2}$$

The law of large numbers ensures that the empirical mean converges almost surely to the expected value:

$$\bar{\mu}(\boldsymbol{X}) = \frac{1}{N}\sum_{i=1}^{N}\mu(\boldsymbol{X}^i) \xrightarrow[N\to+\infty]{} \mathbb{E}_p\left(\mu(\boldsymbol{X})\right) \tag{3.3}$$

The variance of the Monte Carlo estimator $\bar{\mu}(\boldsymbol{X})$ is equal to:

$$\mathbb{V}\mathrm{ar}_p(\bar{\mu}(\boldsymbol{X})) = \frac{\sigma^2}{N} \tag{3.4}$$

where

$$\sigma^2 = \int \left(\mu(\boldsymbol{X}) - \mathbb{E}_p(\mu(\boldsymbol{X}))\right)^2 p(\boldsymbol{X})\, d\boldsymbol{X} \tag{3.5}$$

The law of the error $\mu(\boldsymbol{X}) - \mathbb{E}_p(\mu(\boldsymbol{X}))$, when the number of samples N tends to infinity, is given by the central limit theorem:

$$\frac{\sqrt{N}}{\sigma}(\mu(\boldsymbol{X}) - \mathbb{E}_p(\mu(\boldsymbol{X}))) \xrightarrow[N\to+\infty]{} \mathcal{N}(0,1) \tag{3.6}$$

The average error is of the order of $\frac{\sigma}{\sqrt{N}}$. The dimension of the state is not important in the Monte Carlo approximation error. Monte Carlo methods are based on the simulation of random variables to compute approximately high dimensional integrals. They draw their justification from the law of large numbers which allows to approximate a probability measure by the empirical measure computed from a samples set.

3.1.1 *Generation of Random Variables using the Acceptance-Rejection Method*

The Monte Carlo estimator draws a set of samples using directly the target density $p(\boldsymbol{X})$ that is unknown in practice. In order to obtain N i.i.d. samples using the target density p, the rejection (or acceptance-rejection) method [26] is often employed. Assume that p is known up to a constant multiplier and that we have a proposal law q according to which it is possible to obtain samples directly, and a constant l such that for all $\boldsymbol{X}$, $p(\boldsymbol{X}) \leq l\, q(\boldsymbol{X})$. The rejection method consist to sample $\bar{\boldsymbol{X}} \sim q$ and $u \sim \mathcal{U}([0,1])$ then to consider $u \leq \frac{p(\bar{\boldsymbol{X}})}{l\, q(\bar{\boldsymbol{X}})}$: if this inequality is verified, the sample $\bar{\boldsymbol{X}}$ is accepted and $\boldsymbol{X} = \bar{\boldsymbol{X}}$, else the sample is rejected. The resulting random variable admits for law p. The rejection method requires the knowledge of the proposal density q and a constant l, and its computational cost can become significant when the probability of acceptance $\frac{1}{l}$ is small.

3.1.2 *Importance Sampling*

Alternatively, the samples $\boldsymbol{X}^i$ can be drawn from a proposal density q. A weight w^i is associated at each sample $\boldsymbol{X}^i$ such that $\tilde{w}^i = \frac{p(\boldsymbol{X}^i)}{q(\boldsymbol{X}^i)}$ is the unnormalised weight and $w^i = \frac{\tilde{w}^i}{\sum_{j=1}^N \tilde{w}^j}$ the normalised one. The resulting approximation is called the importance sampling estimator and is defined by:

$$\tilde{\mu}(\boldsymbol{X}) \coloneqq \sum_{i=1}^{N} w^i \mu(\boldsymbol{X}^i), \quad \boldsymbol{X}^i \sim q \tag{3.7}$$

In general, the importance sampling estimator is less efficient than the Monte Carlo estimator since the samples are not distributed according to the target density p. The effective sample size criterion [49] measures the loss of efficiency using $\tilde{\mu}$ instead of $\bar{\mu}$ (see Appendix b for more details).

The sequential importance sampling approximations are called PFs.

3.2 ENSEMBLE KALMAN FILTER

In 1994, a method based on Monte Carlo sampling was proposed: the EnKF [28]. In this Monte Carlo method, the empirical state density is approximated by a set of possible state samples. The main difference with the UKF (see Section 2.5) is that the conditional density is approximated by samples chosen in a stochastic way. EnKF was first applied to linear Gaussian systems [28] and then extended to non-linear Gaussian models [29]. For nonlinear cases, the Gaussian density is generally an approximation of the theoretical state density, since non-linear models lead to non-Gaussian densities.

The state evolution is assumed to be a non-linear Gaussian model and the observation model is assumed to be linear and Gaussian. The state-space model is defined by:

$$\begin{cases} \boldsymbol{X}_k = b_k(\boldsymbol{X}_{k-1}) + \boldsymbol{W}_k \\ \boldsymbol{Y}_k = \boldsymbol{H}_k \ \boldsymbol{X}_k + \boldsymbol{V}_k \end{cases} \tag{3.8}$$

The process and observation noises are assumed to be mutually independent zero-mean white Gaussian noises, with covariances $\boldsymbol{Q}_k \in \mathbb{R}^{d\times d}$ and $\boldsymbol{R}_k \in \mathbb{R}^{d_y \times d_y}$ respectively. The initial law of the state $p(\boldsymbol{X}_0)$ is not necessarily Gaussian but is usually a centered Gaussian of covariance $\boldsymbol{P}_0$. The state density is estimated by a N-set of state samples $\boldsymbol{X}_k^1, \ldots, \boldsymbol{X}_k^N$.

The EnKF's equations are presented in Algorithm 3.1. The notation $\mathbf{0}_d$ denotes the zero value vector in $\mathbb{R}^d$.

In the correction step, the measurement is perturbed with N samples drawn from the law of the observation noise. This is necessary to keep the covariance consistent with the classic Gaussian KF [15]. A set of perturbations $\{\boldsymbol{V}_k^i\}_{i\in[1,N]}$ is obtained and each state sample $\boldsymbol{X}_{k|k-1}^i$ is associated with the perturbed measurements $\boldsymbol{Y}_k + \boldsymbol{V}_k^i$.

Algorithm 3.1 Ensemble Kalman Filter

Initialisation: The initial N-set is drawn as $\{\boldsymbol{X}_0^i\}_{i\in[1,N]}$ using the initial density $p(\boldsymbol{X}_0) = \mathcal{N}(\boldsymbol{X}_0; \mathbf{0}_d, \boldsymbol{P}_0)$.
for each time-step k **do**
  Prediction: Propagate the state samples using the dynamical model $\boldsymbol{X}_{k|k-1}^i = b_k(\boldsymbol{X}_{k-1}^i) + \boldsymbol{W}_k^i$ where $\boldsymbol{W}_k^i \sim \mathcal{N}(\boldsymbol{W}_k^i; 0, \boldsymbol{Q}_k)$.
  Compute the sample mean $\boldsymbol{m}_{k|k-1}^N = \frac{1}{N}\sum_{i=1}^N \boldsymbol{X}_{k|k-1}^i$ and the sample covariance

$$\boldsymbol{P}_{k|k-1}^N = \frac{1}{N-1}\sum_{i=1}^N \left(\boldsymbol{X}_{k|k-1}^i - \boldsymbol{m}_{k|k-1}^N\right)\left(\boldsymbol{X}_{k|k-1}^i - \boldsymbol{m}_{k|k-1}^N\right)^\top$$

  Correction:
  $\boldsymbol{X}_k^i = \boldsymbol{X}_{k|k-1}^i + \boldsymbol{K}_k(\boldsymbol{Y}_k - \boldsymbol{H}_k\, \boldsymbol{X}_{k|k-1}^i + \boldsymbol{V}_k^i)$ where $\boldsymbol{V}_k^i \sim \mathcal{N}(\boldsymbol{V}_k^i, 0, \boldsymbol{R}_k)$.
  $\boldsymbol{K}_k$ is the Kalman gain: $\boldsymbol{K}_k = \boldsymbol{P}_{k|k-1}^N \boldsymbol{H}_k^\top \left(\boldsymbol{H}_k \boldsymbol{P}_{k|k-1}^N \boldsymbol{H}_k^\top + \boldsymbol{R}_k\right)^{-1}$
end for
Return the sample mean $\boldsymbol{m}_k^N$ and covariance $\boldsymbol{P}_k^N$ calculated from the corrected state samples $\boldsymbol{X}_k^i$, $\forall k$.

EnKF was also derived for nonlinear observation model [41]. When the observation equation is nonlinear, i.e., $\boldsymbol{Y}_k = h_k(\boldsymbol{X}_k) + \boldsymbol{V}_k$, the correction step is summarised by the following equations:

$$\boldsymbol{X}_k^i = \boldsymbol{X}_{k|k-1}^i + \boldsymbol{K}_k(\boldsymbol{Y}_k - \boldsymbol{Y}_k^i + \boldsymbol{V}_k^i) \tag{3.9a}$$

$$\boldsymbol{Y}_k^i = h_k(\boldsymbol{X}_{k|k-1}^i), \quad \hat{\boldsymbol{Y}}_k = \frac{1}{N}\sum_{i=1}^N \boldsymbol{Y}_k^i \tag{3.9b}$$

$$\boldsymbol{K}_k = \boldsymbol{P}_k^{XY}\left(\boldsymbol{P}_k^Y\right)^{-1} \tag{3.9c}$$

$$\boldsymbol{P}_k^{XY} = \frac{1}{N-1}\sum_{i=1}^N \left(\boldsymbol{X}_{k|k-1}^i - \boldsymbol{m}_{k|k-1}^N\right)\left(\boldsymbol{Y}_k^i - \hat{\boldsymbol{Y}}_k\right)^\top \tag{3.9d}$$

$$\boldsymbol{P}_k^Y = \frac{1}{N-1}\sum_{i=1}^N \left(\boldsymbol{Y}_k^i - \hat{\boldsymbol{Y}}_k\right)\left(\boldsymbol{Y}_k^i - \hat{\boldsymbol{Y}}_k\right)^\top + \boldsymbol{R}_k \tag{3.9e}$$

The prediction step remains unchanged.

In the EnKF presented in Algorithm 3.1, the perturbation of the measurements introduces an additional sampling noise that can lower the filter performance. Furthermore, the Kalman gain formulation leads to the computationally expensive inversion of the matrix $\boldsymbol{H}_k \boldsymbol{P}_{k|k-1}^N \boldsymbol{H}_k^\top + \boldsymbol{R}_k$. These two drawbacks have motivated the development of deterministic approaches called ensemble square-root filters (see e.g., [3, 11, 85]). These methods avoid sampling issues by

generating an ensemble with the desired sample mean and covariance. The correction step in Algorithm 3.1 is replaced by the following equation:

$$\boldsymbol{X}_k^i = \sum_{j=1}^{N} \left(\boldsymbol{X}_{k|k-1}^j - \boldsymbol{m}_{k|k-1}^N \right) \boldsymbol{W}^{j,i} + \boldsymbol{m}_{k|k-1}^N \tag{3.10}$$

The matrix $\boldsymbol{W}^{j,i}$ is symmetric and depends on the ensemble covariance, the difference between the predicted ensembles and the measurements, and the observation covariance. These methods often require inflation and localisation techniques as sampling errors can lead to instability. Covariance inflation [4, 5] consists in artificially increasing the variance of each state vectors components to correct systematic underestimation of the ensemble covariance. Such a deficiency yields an over confidence on the model dynamics and hence reduces the impact of the measurements.

The main difference between EnKF and PFs lies in the correction phase. In the EnKF, the samples are redirected, while in the PF they are weighted. As a result, each methods has a specific advantage towards the other:

- EnKF samples are redistributed in significant areas of the density during the correction step, which reduces the number of required particles.
- PF is based on an exact probabilistic formulation of the conditional density even in the case of non-linear dynamics.

3.3 PARTICLE FILTERS

Consider the state-space model given in Chapter 2 by equations (2.17) and (2.2):

$$\begin{cases} \boldsymbol{X}_k = b_k(\boldsymbol{X}_{k-1}) + \boldsymbol{W}_k \\ \boldsymbol{Y}_k = h_k(\boldsymbol{X}_k) + \boldsymbol{V}_k \end{cases} \tag{3.11}$$

where $b_k : \mathbb{R}^d \to \mathbb{R}^d$ and $h_k : \mathbb{R}^d \to \mathbb{R}^{d_y}$ are respectively the dynamical and observation functions, $\boldsymbol{W}_k$ and $\boldsymbol{V}_k$ are the process and observation white noises, and with the convention $\boldsymbol{Y}_0 = 0$. The initial state $\boldsymbol{X}_0$ of known density $p(\boldsymbol{X}_0)$ is independent from both noises. The initial state density $p(\boldsymbol{X}_0)$ quantifies the initial state uncertainty. The observations $\{\boldsymbol{Y}_k\}_{k\geq 0}$ are mutually independent conditionally to the state $\{\boldsymbol{X}_k\}_{k\geq 0}$. The noises $\boldsymbol{W}_k$ and $\boldsymbol{V}_k$ are assumed mutually independent. When the dynamical and observation functions are not linear and the noises are not Gaussian, more general estimation methods were proposed as particle filters.

3.3.1 *Sequential Importance Resampling Particle Filter*

The Sequential Importance Resampling Particle Filter (SIR-PF) approach [34, 47, 48] was introduced as an heuristic genetic type algorithm to tackle non-linear non-Gaussian dynamics

and/or severely non-linear measurements. See Del Moral [20] for the first rigorous and well founded article on particle filters. PF consists of empirically estimating the conditional density by a set of weighted points $\left\{\boldsymbol{X}_k^i \in \mathbb{R}^d,\ w_k^i \in \mathbb{R}^+\right\}$ called particles. The particles are a set of i.i.d. samples. They are associated with positive weights w_k^i whose sum is equal to 1. The estimated conditional density is thus defined by a mixture of weighted Dirac deltas, as illustrated in Figure 3.1.

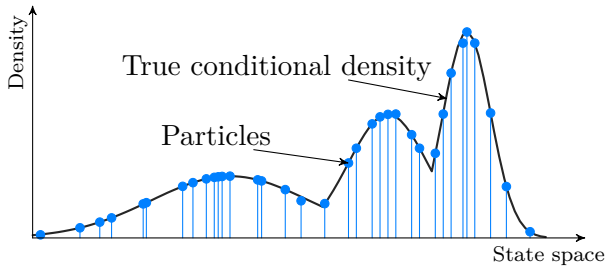

Figure 3.1: Particle filter scheme: the true conditional density is approximated with a mixture of Dirac deltas. Each particle is represented by a point.

The estimated conditional density is defined by:

$$\hat{p}(\boldsymbol{X}_k|\mathcal{Y}_k) = \sum_{i=1}^{N} w_k^i \delta(\boldsymbol{X}_k - \boldsymbol{X}_k^i) \tag{3.12}$$

where δ is the Dirac delta function on $\mathbb{R}^d$ equal to 0 on $\mathbb{R}^d \backslash \{\mathbf{0}_d\}$ and to infinity on $\mathbf{0}_d$. The predicted particles $\mathcal{X}_k^i$ are drawn from a suitable proposal density $q(\mathcal{X}_k|\mathcal{Y}_k)$. The proposal density is chosen to factorize such that:

$$q(\mathcal{X}_k|\mathcal{Y}_k) \coloneqq q(\boldsymbol{X}_k|\mathcal{X}_{k-1}, \mathcal{Y}_k)\ q(\mathcal{X}_{k-1}|\mathcal{Y}_{k-1}) \tag{3.13}$$

Furthermore, we consider that the proposal density depends only on the previous state vector and the current measurement, such that:

$$q(\boldsymbol{X}_k|\mathcal{X}_{k-1}, \mathcal{Y}_k) = q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k) \tag{3.14}$$

The choice of the proposal density will be discussed afterwards.

The weights are updated according to the likelihood $p(\boldsymbol{Y}_k|\boldsymbol{X}_k)$, the transition density $p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1})$, and the proposal density $q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)$:

$$w_k^i \propto w_{k-1}^i \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k^i)\ p(\boldsymbol{X}_k^i|\boldsymbol{X}_{k-1}^i)}{q(\boldsymbol{X}_k^i|\boldsymbol{X}_{k-1}^i, \boldsymbol{Y}_k)} \tag{3.15}$$

A detailed derivation of the weights update equation (3.15) can be find in Ristic, Arulampalam, and Gordon [76].

When the number of particles N tends toward infinity, the approximation of the conditional density (3.12) approaches the true conditional density $p(\boldsymbol{X}_k|\mathcal{Y}_k)$. Formulation (3.12) provides an estimation of the conditional density. For practical use, one can derive a point-wise state estimate. The Maximum A Posteriori (MAP) (2.6) can be numerically computed but MAP computation often yields a high computational load [74]. In practice, the estimate is often computed via a Least Square (LS) approximation (2.7), which corresponds to the barycentre of the weighted particle cloud, under the assumption that the actual state density tends to a Gaussian density. The state estimate is then computed by:

$$\hat{\boldsymbol{X}}_k = \sum_{i=1}^{N} w_k^i \ \boldsymbol{X}_k^i \tag{3.16}$$

and its covariance by:

$$\boldsymbol{P}_k = \sum_{i=1}^{N} w_k^i \ (\boldsymbol{X}_k^i - \hat{\boldsymbol{X}}_k)(\boldsymbol{X}_k^i - \hat{\boldsymbol{X}}_k)^\top \tag{3.17}$$

Degeneracy Problem

For proposal densities of the form (3.14), it was shown that the variance of the weights can only increase over time [27]. The increasing variance has a harmful effect on the accuracy of the estimate and leads to a common problem: the degeneracy problem. After a certain number of recursive steps, all but one particle will have negligible weights. Since the weights are normalised, one weight will tend to one while all the other weights tend to zero. When the proposal density differs from the conditional density, degeneracy phenomenon raises. A large computational effort is devoted to updating particles for which the contribution to the approximation of the conditional density is almost zero.

In order to keep a representative set of particles, two methods (that can be combined) were proposed. A first solution is to add a resampling step [34] which duplicates the high-weighted particles to replace low-weighted particles while keeping the total number of particles constant. A second solution consists in choosing an appropriate proposal density which minimises the variance of the weights.

Resampling

The resampling step consists of replacing the current set of particles with a new one which described the state density in a more accurate way. The particles with high normalised weights are selected and low-weighted particles are discarded. The selected particles are duplicated in order to keep a constant total number of particles N.

A large variety of sampling and resampling algorithms can be used in the context of PF (see [23] and [54] for a review on resampling methods). A resampling algorithm provides an integer value indicating how many duplications of each particle will occur in the new sample: 0 if the particle is discarded, 1 if it is kept, and n if it is duplicated into n instances.

The multinomial resampling is commonly used and a brief description of this method is provided here. This algorithm defines the number of duplications of each particle with regard to the value of the weights. It is achieved by drawing N independent values $\{u^i\}_{i\in[1,N]}$ following a uniform law on $[0,1]$ and comparing them to the cumulative sum of the weights. A duplication number n^i can be computed for each particle, corresponding to the number of uniform values lying into their cumulative interval. The multinomial resampling is summarised in Algorithm 3.2. After the resampling step, weights are reset to $\frac{1}{N}$.

Algorithm 3.2 Multinomial Resampling

1: **Initialisation:** Set the duplication counters $n^i = 0$, $\forall i \in [1, N]$.
2: **for each** $i \in [1, N]$ **do**
3: | Draw $u^i \sim \mathcal{U}_{[0,1]}$
4: | Find $j \in [1, N]$ such that $u^i \in \left]\sum_{l=1}^{j-1} w^l, \sum_{l=1}^{j} w^l\right[$ where $\{w^i\}_{i\in[1,N]}$ are the particle weights.
5: | Count $n^j = n^j + 1$
6: **end for**
7: Return $\{n^i\}_{i\in[1,N]}$.

The resampling step reduces the effects of degeneracy, but it introduces practical and theoretical issues. From an implementation point of view, the possibility of parallelizing the multinomial resampling is reduced since all the particles must be combined. On the theoretical level, the resampling introduces a dependency between the particles, making the convergence results difficult to establish.

The particles that have high weights are statistically selected many times. This leads to a loss of diversity among the particles. This problem is known as sample impoverishment. Because of these drawbacks, the resampling step is performed only when necessary. Several triggering criteria have been proposed to detect degeneracy. The most commonly used is the effective sample size [49] defined by:

$$\mathrm{ESS} = \frac{N}{1 + \mathbb{V}\mathrm{ar}_q(\tilde{w}_k)} \tag{3.18}$$

where $\{\tilde{w}_k\}$ are the unnormalised weights. The effective sample size ESS cannot be evaluated exactly. An estimate is given by (see Appendix b for details):

$$\widehat{\mathrm{ESS}} = \frac{1}{\sum_{i=1}^{N} \left(w_k^i\right)^2} \tag{3.19}$$

where $1 \leq \widehat{\mathrm{ESS}} \leq N$. The approximate effective sample size $\widehat{\mathrm{ESS}}$ is equal to the number of samples N, when all the weights are set to $\frac{1}{N}$. When this criterion is small, it indicates severe degeneracy. The resampling is triggered whenever $\widehat{\mathrm{ESS}} < N_{\mathrm{th}}$. The threshold N_{th} is usually equal to $\theta_{\mathrm{th}} N$ where $\theta_{\mathrm{th}} \in [0,1]$ is a tuning parameter. Other degeneracy criteria exist such as the entropy [71] which measures the dispersion of the particle weights.

Choice of the Importance Density

The choice of the proposal density is one of the most critical issues in the design of a PF. Ideally, the proposal density function should be the unknown conditional density $p(\boldsymbol{X}_k|\mathcal{Y}_k)$. The so-called optimal proposal density function that minimises the variance of the local weights, conditioned upon $\boldsymbol{X}_{k-1}^i$ and $\boldsymbol{Y}_k$, was shown to be $q_{opt}(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}^i, \boldsymbol{Y}_k) = p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)$ [21, 27]. We emphasize that this proposal density is not globally optimal. The analytic evaluation of this density is difficult for most cases. The optimal proposal density can be used in two specific situations: when $\boldsymbol{X}_k$ is a member of a finite set or when the density $p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)$ is Gaussian (see [27] or example 3 in Del Moral [21]).

It is possible to construct suboptimal approximations to the optimal proposal density by using, e.g., local linearisation techniques [27] or Monte Carlo approximations (see Section 4.2.2 in Del Moral [21]). Local linearisation techniques use a proposal density that is a Gaussian approximation of $p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)$ and are based on EKF or UKF [27, 86].

To facilitate the implementation of the filter, the proposal density is often taken equal to the transition density $q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k) = p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1})$. Filters using this choice of proposal density are called bootstrap filters [34]. The weights update (3.15) is therefore simplified by:

$$w_k^i \propto w_{k-1}^i \, p(\boldsymbol{Y}_k|\boldsymbol{X}_k^i) \tag{3.20}$$

When the transition density is used as the proposal density, the particle weights can degenerate rapidly if either of the following situations occur:

- When the prior density is not very informative and the likelihood has very pronounced modes, the correction is ineffective as many particles have weights close to zero and thus do not participate in the approximation of the conditional density. This situation arises when the observation noise is low. This case is illustrated in Figure 3.2a.
- The Figure 3.2b illustrates the case of inconsistency between the likelihood and the prior density (i.e., the support of the prior density does not contain the support of the likelihood density). This situation arises when the process noise is too low. It is possible that few particles are placed in high likelihood regions as the propagation of the particles do not consider the current measurement.

Both situations can occur in sensors-based navigation applications. The use of very precise sensors may lead to a very sharp likelihood, which favors the appearance of the first situation. The inconsistency situation can occur when the system has a high confidence in an imprecise dynamic model (i.e., when the process noise is too low).

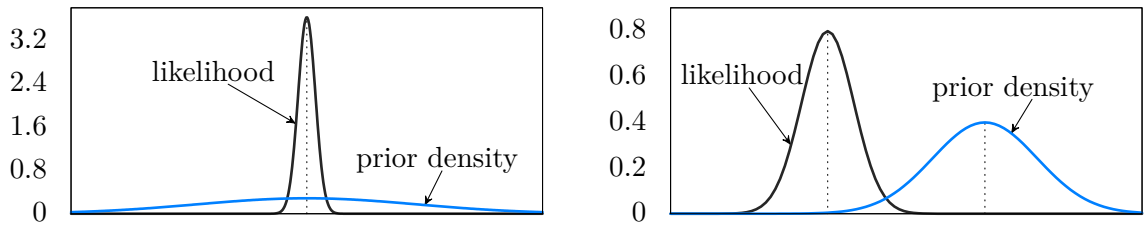

(a) Sharp likelihood and non-informative prior density.

(b) Inconsistency between the likelihood and the prior density.

Figure 3.2: Sharp likelihood and inconsistency schemes.

Several methods were introduced to replace the particles in high likelihood regions, e.g., the auxiliary particle filter [72] that incorporates the current measurement in the prior density, or the progressive correction [68] that introduces intermediate distributions between the prior and the likelihood densities.

In the three filters presented hereafter (namely the SIR-PF, RPF, and RBPF), the proposal density is equal to the transition density.

Sequential Importance Resampling Particle Filter Algorithm

Particle filters using a resampling step belong to the SIR-PF category. The SIR-PF equations are summarised in Algorithm 3.3.

Algorithm 3.3 Sequential Importance Resampling Particle Filter

Initialisation: The initial particle set is drawn as $\{\boldsymbol{X}_0^i\}_{i\in[1,N]}$ using the initial density $p(\boldsymbol{X}_0)$ and the initial weights set $\{w_0^i\}_{i\in[1,N]}$ is taken equal to $\frac{1}{N}$.
for each time-step k **do**
  Prediction: Draw the particles from the transition density $\boldsymbol{X}_k^i \sim p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}^i)$.
  Correction: Update the weights $w_k^i \propto w_{k-1}^i\, p(\boldsymbol{Y}_k|\boldsymbol{X}_k^i)$.
  Compute the state estimate $\hat{\boldsymbol{X}}_k$ (3.16) and its covariance $\boldsymbol{P}_k$ (3.17).
  if a resampling criterion is satisfied, e.g., $\widehat{\mathrm{ESS}} < N_{\mathrm{th}}$, see (3.19) **then**
    Draw a new set of particles $\{\boldsymbol{X}_k^i\}_{i\in[1,N]}$ using a resampling method, e.g., the multinomial resampling (Algorithm 3.2).
    Reset the weights to $\frac{1}{N}$.
  end if
end for
Return the state estimate $\hat{\boldsymbol{X}}_k$ and its covariance $\boldsymbol{P}_k$, $\forall k$.

SIR-PF is able to tackle a larger variety of problems than parametric filters. The resampling was introduced in order to tackle the degeneracy problem. However, as discussed in the

previous paragraphs, sample impoverishment may occur especially when the proposal density is taken equal to the transition density. When the observation noise is too low (see Figure 3.2a), the actual conditional density may concentrate in an area lacking of particles, which leads to ill estimated weights. When the process noise is too low (see Figure 3.2b), the resampling tends to select only a few particles and will asymptotically lead to N copies of a single one. The estimate is thus not guaranteed to remain close to the actual state. Some derivations of SIR-PF were introduced to make it more stable to low noise systems. RPF [61] proposed an additional step of particle perturbations during the resampling step. The purpose of the regularisation is to ensure that the assumption of samples independence is respected by injecting a small noise in the signal. RPF is based on the kernel estimation theory [81], and replaces the discrete approximation of the conditional density by a continuous approximation.

3.3.2 *Regularised Particle Filter*

RPF [60, 61, 67] replaces the conditional density approximation given by the mixture of weighted Dirac deltas (3.12), by a mixture of weighted kernels centered on each particle $\boldsymbol{X}_k^i$:

$$\hat{p}(\boldsymbol{X}_k|\mathcal{Y}_k) = \sum_{i=1}^{N} w_k^i \, \mathcal{K}_h(\boldsymbol{X}_k - \boldsymbol{X}_k^i) \tag{3.21}$$

where the kernel $\mathcal{K}_h$ is defined by:

$$\mathcal{K}_h(\boldsymbol{X}) \coloneqq \frac{1}{h^d} \mathcal{K}\left(\frac{1}{h}\boldsymbol{X}\right) \tag{3.22}$$

with $\mathcal{K}$ a reference kernel from $\mathbb{R}^d$ to $\mathbb{R}$ which integrates to unity on $\mathbb{R}^d$ and h the associated bandwidth. The kernel is assumed to be symmetric such that $\mathcal{K}_h(-\boldsymbol{X}) = \mathcal{K}_h(\boldsymbol{X})$. The approximation of the conditional density is illustrated in Figure 3.3.

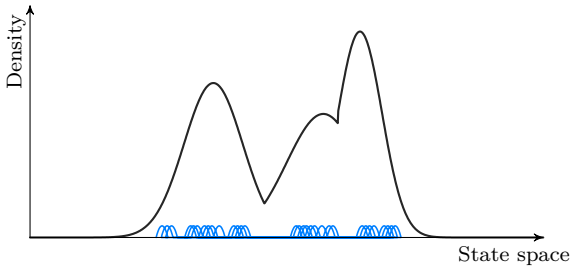

Figure 3.3: Kernel estimator used to estimate the conditional density from a set samples $\{\boldsymbol{X}^i\}_{i\in[1,N]}$.

The kernel and the bandwidth are chosen to minimise the Mean Integrated Square Error (MISE) between the true conditional density $p(\boldsymbol{X}_k|\mathcal{Y}_k)$ and the corresponding regularised empirical density $\hat{p}(\boldsymbol{X}_k|\mathcal{Y}_k)$ (3.21), which is defined as:

$$\text{MISE}(\hat{p}, p) = \mathbb{E}\left[\int_{\mathbb{R}^d} \left(\hat{p}(\boldsymbol{X}_k|\mathcal{Y}_k) - p(\boldsymbol{X}_k|\mathcal{Y}_k)\right)^2 d\boldsymbol{X}_k\right] \tag{3.23}$$

When all the particles have the same weight, which is the case immediately after the resampling step, the optimal kernel, which minimises the MISE, is the Epanechnikov kernel:

$$\mathcal{K}_{\text{opt}}(\boldsymbol{X}) = \begin{cases} \frac{d+2}{2\,c_d}\left(1 - \|\boldsymbol{X}\|^2\right) & \text{if } \|\boldsymbol{X}\| < 1 \\ 0 & \text{otherwise} \end{cases} \tag{3.24}$$

where c_d is the volume of the unit hypersphere in $\mathbb{R}^d$. The associated optimal bandwidth is:

$$h_{\text{opt}} = \mu\, A(\mathcal{K})\, N^{-\frac{1}{d+4}} \tag{3.25}$$

where

$$A(\mathcal{K}) = \left[8\, c_d^{-1}\,(d+4)(2\sqrt{\pi})^d\right]^{\frac{1}{d+4}} \tag{3.26}$$

$0 < \mu < 1$ is a tuning parameter introduced to limit the impact of the regularisation on each individual mode, when the assumption of unimodality of the conditional density is not satisfied.

Algorithm 3.4 Regularised Resampling Procedure

1: **if** a resampling criterion is satisfied, e.g., $\widehat{\text{ESS}} < N_{\text{th}}$, see (3.19) **then**
2: Draw a new set of particles $\{\boldsymbol{X}_k^i\}_{i\in[1,N]}$ using a resampling method, e.g., the multinomial resampling (Algorithm 3.2).
3: [Regularisation:] Compute the Cholesky decomposition of the empirical covariance of the resampled set of particles $\boldsymbol{A}_k$ such that:

$$\boldsymbol{P}_k = \boldsymbol{A}_k\, \boldsymbol{A}_k^\top = \sum_{i=1}^{N} w_k^i \left(\boldsymbol{X}_k^i - \hat{\boldsymbol{X}}_k\right)\left(\boldsymbol{X}_k^i - \hat{\boldsymbol{X}}_k\right)^\top \tag{3.27}$$

4: Draw $\{\kappa_k^i\}_{i\in[1,N]}$ samples from the regularisation kernel $\mathcal{K}_{\text{opt}}$ (3.24).
5: The regularised particles are obtained by adding a noise to each particle:

$$\boldsymbol{X}_k^i = \boldsymbol{X}_k^i + h_{\text{opt}}\ \boldsymbol{A}_k\ \kappa_k^i \tag{3.28}$$

where h_{opt} is the optimal bandwidth (3.25).
6: Reset the weights to $\frac{1}{N}$.
7: **end if**
8: Return the new set of particles $\{\boldsymbol{X}_k^i\}_{i\in[1,N]}$ and the weights $w_k^i = \frac{1}{N}$, $\forall i$.

The optimal bandwidth (3.25) is expressed for a normalised density whose covariance is a d dimensional identity matrix. In practice, the kernel bandwidth must be adapted to the density's covariance. This can be done by computing the Cholesky decomposition of the empirical covariance of the resampled set of particles. RPF algorithm differs from the generic SIR-PF (see Algorithm 3.3) only in its resampling step. The regularised resampling procedure is summarised in Algorithm 3.4.

The regularisation step improves the sample diversity by adjusting the empirical density. In practical scenarios, RPF performs better than SIR-PF in cases of sample impoverishment (e.g., when the process noise is low) [61].

The kernel estimation leads a variety of kernel-based particle filters, e.g., the Kalman-particle kernel filter (KPKF [19, 70]), or the box regularised particle filter (BRPF, [59]).

3.3.3 *Rao-Blackwellised Particle Filter*

Although Monte Carlo methods are in principle applicable in any dimension, the number of required particles N needs to be high for higher-dimensional systems [37]. For some classes of state-space models the number of particles in the particle filter can be reduced using a variance reduction method known as Rao-Blackwellisation [17, 64, 80]. The RBPF (also called Marginalised Particle Filter) results in a higher robustness, since particle Dirac deltas are maginalised on the non-linear dimensions only, which limits the curse of dimensionality [32].

The RBPF is an efficient implementation for conditionally linear Gaussian models. Assuming that the state vector $\boldsymbol{X}_k$ can be partitioned into two sub-vectors $\boldsymbol{X}_k = \left[\boldsymbol{X}_k^n, \boldsymbol{X}_k^l\right]^\top$, where $\boldsymbol{X}_k^n \in \mathbb{R}^n$ denotes the nonlinear state variable and $\boldsymbol{X}_k^l \in \mathbb{R}^l$ denotes the state variable with conditionally linear dynamics (such that $d = n + l$), the state-space model takes the following form:

$$\begin{cases} \boldsymbol{X}_{k+1}^n = b_k^n(\boldsymbol{X}_k^n) + \boldsymbol{B}_k^n(\boldsymbol{X}_k^n)\,\boldsymbol{X}_k^l + \boldsymbol{W}_k^n \\ \boldsymbol{X}_{k+1}^l = b_k^l(\boldsymbol{X}_k^n) + \boldsymbol{B}_k^l(\boldsymbol{X}_k^n)\,\boldsymbol{X}_k^l + \boldsymbol{W}_k^l \\ \boldsymbol{Y}_k = h_k(\boldsymbol{X}_k^n) + \boldsymbol{H}_k(\boldsymbol{X}_k^n)\,\boldsymbol{X}_k^l + \boldsymbol{V}_k \end{cases} \tag{3.29}$$

where the measurement noise $\boldsymbol{V}_k$ is assumed to be a centered Gaussian white noise with covariance $\boldsymbol{R}_k$. The process noise is also assumed white and Gaussian:

$$\boldsymbol{W}_k = \begin{bmatrix} \boldsymbol{W}_k^n \\ \boldsymbol{W}_k^l \end{bmatrix} \sim \mathcal{N}(0_d, \boldsymbol{Q}_k) \tag{3.30}$$

with covariance $\boldsymbol{Q}_k \in \mathbb{R}^{d\times d}$ such that:

$$\boldsymbol{Q}_k = \begin{pmatrix} \boldsymbol{Q}_k^n & \boldsymbol{Q}_k^{ln} \\ (\boldsymbol{Q}_k^{ln})^\top & \boldsymbol{Q}_k^l \end{pmatrix} \tag{3.31}$$

The known initial laws are $p(\boldsymbol{X}_0^n)$ for the nonlinear part and $\mathcal{N}(\boldsymbol{X}_k^l;\, \hat{\boldsymbol{X}}_0^l, \boldsymbol{P}_0^l)$ for the linear part.

The aim is to estimate the conditional density $p(\boldsymbol{X}_k|\mathcal{Y}_k) = p(\boldsymbol{X}_k^n, \boldsymbol{X}_k^l|\mathcal{Y}_k)$. Using the Bayes' rule, the density $p(\mathcal{X}_k^n, \boldsymbol{X}_k^l|\mathcal{Y}_k)$ can be factorised into two parts:

$$p(\mathcal{X}_k^n, \boldsymbol{X}_k^l|\mathcal{Y}_k) = p(\boldsymbol{X}_k^l|\mathcal{X}_k^n, \mathcal{Y}_k)\, p(\mathcal{X}_k^n|\mathcal{Y}_k) \tag{3.32}$$

The density $p(\boldsymbol{X}_k^l|\mathcal{X}_k^n, \mathcal{Y}_k)$ is analytically tractable, as $p(\boldsymbol{X}_k^l|\mathcal{X}_k^n, \mathcal{Y}_k) = \mathcal{N}(\boldsymbol{X}_k^l;\, \hat{\boldsymbol{X}}_{k|k}^l, \boldsymbol{P}_{k|k}^l)$, where the equations for the mean $\hat{\boldsymbol{X}}_{k|k}^l$ and covariance $\boldsymbol{P}_{k|k}^l$ are given by the KF (see Section 2.4). The density $p(\mathcal{X}_k^n|\mathcal{Y}_k)$ can be estimated by a PF:

$$p(\mathcal{X}_k^n|\mathcal{Y}_k) \simeq \sum_{i=1}^{N} w_k^i \delta(\mathcal{X}_k^n - \mathcal{X}_k^{n,i}) \tag{3.33}$$

The marginal density $p(\mathcal{X}_k^l|\mathcal{Y}_k)$ can be approximated by:

$$p(\mathcal{X}_k^l|\mathcal{Y}_k) = \int p(\mathcal{X}_k^l|\mathcal{X}_k^n, \mathcal{Y}_k)\, p(\mathcal{X}_k^n|\mathcal{Y}_k)\; d\mathcal{X}_k^n \tag{3.34a}$$

$$\simeq \int p(\mathcal{X}_k^l|\mathcal{X}_k^n, \mathcal{Y}_k) \sum_{i=1}^{N} w_k^i \delta(\mathcal{X}_k^n - \mathcal{X}_k^{n,i}) d\mathcal{X}_k^n \tag{3.34b}$$

$$\simeq \sum_{i=1}^{N} w_k^i\; p(\mathcal{X}_k^l|\mathcal{X}_k^{n,i}, \mathcal{Y}_k) \tag{3.34c}$$

Similary,

$$p(\boldsymbol{X}_k^l|\mathcal{Y}_k) \simeq \sum_{i=1}^{N} w_k^i\; p(\boldsymbol{X}_k^l|\mathcal{X}_k^{n,i}, \mathcal{Y}_k) \tag{3.35}$$

For the sake of brevity, the dependence of $\boldsymbol{X}_k^n$ in $\boldsymbol{B}_k^n$, $\boldsymbol{B}_k^l$, and $\boldsymbol{H}_k$ are suppressed below. The linear state variables are estimated by the following equations:

$$p(\boldsymbol{X}_k^l|\mathcal{X}_k^n, \mathcal{Y}_k) = \mathcal{N}\left(\boldsymbol{X}_k^l;\, \hat{\boldsymbol{X}}_{k|k}^l, \boldsymbol{P}_{k|k}^l\right) \tag{3.36a}$$

$$p(\boldsymbol{X}_{k+1}^l|\mathcal{X}_{k+1}^n, \mathcal{Y}_k) = \mathcal{N}\left(\boldsymbol{X}_{k+1}^l;\, \hat{\boldsymbol{X}}_{k+1|k}^l, \boldsymbol{P}_{k+1|k}^l\right) \tag{3.36b}$$

where

$$\hat{\boldsymbol{X}}_{k|k}^l = \hat{\boldsymbol{X}}_{k|k-1}^l + \boldsymbol{K}_k \left(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k^n) - \boldsymbol{H}_k\; \hat{\boldsymbol{X}}_{k|k-1}^l\right) \tag{3.37a}$$

$$\boldsymbol{P}_{k|k}^l = \boldsymbol{P}_{k|k-1}^l - \boldsymbol{K}_k \boldsymbol{M}_k \boldsymbol{K}_k^\top \tag{3.37b}$$

$$\boldsymbol{K}_k = \boldsymbol{P}_{k|k-1}^l \boldsymbol{H}_k^\top \boldsymbol{M}_k^{-1} \tag{3.37c}$$

$$\boldsymbol{M}_k = \boldsymbol{R}_k + \boldsymbol{H}_k \boldsymbol{P}_{k|k-1}^l \boldsymbol{H}_k^\top \tag{3.37d}$$

and

$$\hat{\boldsymbol{X}}_{k+1|k}^l = \bar{\boldsymbol{B}}_k^l\; \hat{\boldsymbol{X}}_{k|k}^l + \left(\boldsymbol{Q}_k^{ln}\right)^\top (\boldsymbol{Q}_k^n)^{-1}\; \boldsymbol{Z}_k + b_k^l(\boldsymbol{X}_k^n) + \boldsymbol{L}_k \left(\boldsymbol{Z}_k - \boldsymbol{B}_k^n\; \hat{\boldsymbol{X}}_{k|k}^l\right) \tag{3.38a}$$

$$\boldsymbol{P}_{k+1|k}^l = \bar{\boldsymbol{B}}_k^l \boldsymbol{P}_{k|k}^l \left(\bar{\boldsymbol{B}}_k^l\right)^\top + \bar{\boldsymbol{Q}}_k^l - \boldsymbol{L}_k \boldsymbol{N}_k \boldsymbol{L}_k^\top \tag{3.38b}$$

$$\boldsymbol{N}_k = \boldsymbol{B}_k^n \boldsymbol{P}_{k|k}^l \left(\boldsymbol{B}_k^n\right)^\top + \boldsymbol{Q}_k^n \tag{3.38c}$$

$$\boldsymbol{L}_k = \bar{\boldsymbol{B}}_k^l \boldsymbol{P}_{k|k}^l \left(\boldsymbol{B}_k^n\right)^\top \boldsymbol{N}_k^{-1} \tag{3.38d}$$

where $\boldsymbol{Z}_k = \boldsymbol{X}^n_{k+1} - b^n_k(\boldsymbol{X}^n_k)$.

Covariances $\bar{\boldsymbol{B}}^l_k$ and $\bar{\boldsymbol{Q}}^l_k$ in (3.38) are defined by:

$$\bar{\boldsymbol{B}}^l_k = \boldsymbol{B}^l_k - \left(\boldsymbol{Q}^{ln}_k\right)^\top (\boldsymbol{Q}^n_k)^{-1} \boldsymbol{B}^n_k \tag{3.39a}$$

$$\bar{\boldsymbol{Q}}^l_k = \boldsymbol{Q}^l_k - \left(\boldsymbol{Q}^{ln}_k\right)^\top (\boldsymbol{Q}^n_k)^{-1} \boldsymbol{Q}^{ln}_k \tag{3.39b}$$

The density $p(\mathcal{X}^n_k|\mathcal{Y}_k)$ (3.32) will be approximated using the standard PF equations. From the Bayes' theorem and the Markov property, this density can be write as:

$$p(\mathcal{X}^n_k|\mathcal{Y}_k) = \frac{p(\boldsymbol{Y}_k|\mathcal{X}^n_k, \mathcal{Y}_{k-1})\; p(\boldsymbol{X}^n_k|\mathcal{X}^n_{k-1}, \mathcal{Y}_{k-1})}{p(\boldsymbol{Y}_k|\mathcal{Y}_{k-1})}\, p(\mathcal{X}^n_{k-1}|\mathcal{Y}_{k-1}) \tag{3.40}$$

where an approximation of $p(\mathcal{X}^n_{k-1}|\mathcal{Y}_{k-1})$ is given by the previous iteration of the PF. The analytic expressions for $p(\boldsymbol{Y}_k|\mathcal{X}^n_k, \mathcal{Y}_{k-1})$ and $p(\boldsymbol{X}^n_k|\mathcal{X}^n_{k-1}, \mathcal{Y}_{k-1})$ are given by:

$$p(\boldsymbol{Y}_k|\mathcal{X}^n_k, \mathcal{Y}_{k-1}) = \mathcal{N}(\boldsymbol{Y}_k;\, h_k(\boldsymbol{X}^n_k) + \boldsymbol{H}_k\, \hat{\boldsymbol{X}}^l_{k|k-1},\, \boldsymbol{H}_k \boldsymbol{P}^l_{k|k-1} \boldsymbol{H}_k^\top + \boldsymbol{R}_k) \tag{3.41a}$$

$$p(\boldsymbol{X}^n_{k+1}|\mathcal{X}^n_k, \mathcal{Y}_k) = \mathcal{N}(\boldsymbol{X}^n_{k+1};\, b^n_k(\boldsymbol{X}^n_k) + \boldsymbol{B}^n_k\, \hat{\boldsymbol{X}}^l_{k|k},\, \boldsymbol{B}^n_k \boldsymbol{P}^l_{k|k} \left(\boldsymbol{B}^n_k\right)^\top + \boldsymbol{Q}^n_k) \tag{3.41b}$$

The details of the derivation of the RBPF equations can be found in Schön [79]. The principle of RBPF is to apply a Kalman filter at each particles of the nonlinear state variables $\boldsymbol{X}^{n,i}_{k+1}$. RBPF is summarised in Algorithm 3.5.

Similarly to SIR-PF, the state estimate (3.16) and its covariance (3.17) can be computed for the nonlinear part of the state $\boldsymbol{X}^n_k$, using the weights w^i_k and the particles $\boldsymbol{X}^{n,i}_k$. Equivalent formulas for the linear part of the state $\boldsymbol{X}^l_k$ based on the mean and covariance samples were derived [64]:

$$\hat{\boldsymbol{X}}^l_k = \sum_{i=1}^N w^i_k\; \boldsymbol{X}^{l,i}_k \tag{3.42a}$$

$$\boldsymbol{P}^l_k = \sum_{i=1}^N w^i_k \left(\boldsymbol{P}^{l,i}_{k|k} + (\boldsymbol{X}^{l,i}_k - \hat{\boldsymbol{X}}^l_k)(\hat{\boldsymbol{X}}^{l,i}_k - \boldsymbol{X}^l_k)^\top\right) \tag{3.42b}$$

In Algorithm 3.5, the particle filter used in the RBPF is the SIR-PF. However, the RBPF can be performed with any particle filter, e.g., the auxiliary particle filter or the RPF.

If the same number of particles are used in the SIR-PF and the RBPF, the later will provide better estimates for two reasons:

- The dimension of $p(\boldsymbol{X}^n_k|\mathcal{Y}_k)$ is smaller than the dimension of $p(\boldsymbol{X}^n_k, \boldsymbol{X}^l_k|\mathcal{Y}_k)$, implying that the particles occupy a lower dimensional space.
- Optimal algorithm (i.e., KF) is used in order to estimate the linear state variables.

Although RBPF needs less particles, it is computationally as demanding as standard particle filters as N covariances $\boldsymbol{P}^{l,i}_{k|k}$ are computed at each time-step. When the matrices $\boldsymbol{B}^n_k$, $\boldsymbol{B}^l_k$, and $\boldsymbol{H}_k$ are independent of the nonlinear state variables $\boldsymbol{X}^n_k$ in the state-space model (3.29), the covariance is only update once for each time-step: $\boldsymbol{P}^{l,i}_{k|k} = \boldsymbol{P}^l_{k|k},\ \forall i$. In this case, the gain with respect to the computational load can truly be substantial.

Algorithm 3.5 Rao-Backwellised Particle Filter

Initialisation: The initial particles set is drawn as $\{\boldsymbol{X}_0^{n,i}\}_{i\in[1,N]}$ using the initial density $p(\boldsymbol{X}_0^n)$ and the initial weights set $\{w_0^i\}_{i\in[1,N]}$ is taken equal to $\frac{1}{N}$. The initial mean and covariance of the linear state variables set is $\left\{\boldsymbol{X}_{0|0}^{l,i}, \boldsymbol{P}_{0|0}^{l,i}\right\} = \{\hat{\boldsymbol{X}}_0^l, P_0^l\}$, $\forall i \in [1, N]$.
for $k = 0, 1, \ldots$ **do**
  Particle prediction: Draw the particles $\boldsymbol{X}_{k+1}^{n,i} \sim p(\boldsymbol{X}_{k+1}^n | \mathcal{X}_k^{n,i}, \mathcal{Y}_k)$ (3.41b).
  Kalman prediction: For each particles $\boldsymbol{X}_{k+1}^{n,i}$, sample the mean $\boldsymbol{X}_{k+1|k}^{l,i}$ and covariance $\boldsymbol{P}_{k+1|k}^{l,i}$ (3.38a)-(3.39b) of the density $p(\boldsymbol{X}_{k+1}^l | \mathcal{X}_{k+1}^n, \mathcal{Y}_k)$ (3.36b).
  Particle correction: Update the weights $w_{k+1}^i \propto w_k^i \, p(\boldsymbol{Y}_{k+1} | \mathcal{X}_{k+1}^{n,i}, \mathcal{Y}_k)$ (3.41a).
  if a resampling criterion is satisfied, e.g., $\widehat{\mathrm{ESS}} < N_{\mathrm{th}}$, see (3.19) **then**
    Draw a new set of particles $\{\boldsymbol{X}_{k+1}^i = [\boldsymbol{X}_{k+1}^{n,i}, \boldsymbol{X}_{k+1|k}^{l,i}]^\top\}_{i\in[1,N]}$ using a resampling method, e.g., the multinomial resampling (Algorithm 3.2).
    Reset the weights to $\frac{1}{N}$.
  end if
  Kalman correction: Update the mean $\boldsymbol{X}_{k+1|k+1}^{l,i}$ and covariance $\boldsymbol{P}_{k+1|k+1}^{l,i}$ (3.37a)-(3.37d) of the density $p(\boldsymbol{X}_{k+1}^l | \mathcal{X}_{k+1}^n, \mathcal{Y}_{k+1})$ (3.36a).
  Compute the nonlinear state estimate $\hat{\boldsymbol{X}}_{k+1}^n$ (3.16) and its covariance $\boldsymbol{P}_{k+1}^n$ (3.17).
  Compute the linear state estimate $\hat{\boldsymbol{X}}_{k+1}^l$ (3.42a) and its covariance $\boldsymbol{P}_{k+1}^l$ (3.42b).
end for
Return the state estimate $\hat{\boldsymbol{X}}_{k+1} = [\hat{\boldsymbol{X}}_{k+1}^n, \hat{\boldsymbol{X}}_{k+1}^l]^\top$ and its covariance $\boldsymbol{P}_{k+1} = \begin{pmatrix} \boldsymbol{P}_{k+1}^n & 0_{n\times l} \\ 0_{l\times n} & \boldsymbol{P}_{k+1}^l \end{pmatrix}$, $\forall k$.

3.3.4 *Weighted Ensemble Kalman Filter*

WEnKF [69] consists in using the particle propagation of the EnKF together with a weighting based on the principle of importance sampling presents in particle filters. By doing so, WEnKF combines the main advantages of both filters: the number of particles is reduced due to the efficient EnKF correction and the conditional density is approximated by an exact probabilistic formulation. Indeed, it has been shown in geophysical sciences that for the same performance, WEnKF needs fewer samples than bootstrap particle filters [51, 71]. WEnKF was derived for the two EnKFs variants identified in Section 3.2. Only WEnKF based on EnKF with observation perturbations is presented here and is summarised in Algorithm 3.6.

Algorithm 3.6 Weighted Ensemble Kalman Filter

Initialisation: The initial particles set is drawn as $\{\boldsymbol{X}_0^i\}_{i\in[1,N]}$ using the initial density $p(\boldsymbol{X}_0) \sim \mathcal{N}(\boldsymbol{X}_0; 0_d, \boldsymbol{P}_0)$ and the initial weights set $\{w_0^i\}_{i\in[1,N]}$ is taken equal to $\frac{1}{N}$.
for each time-step k **do**
  Particle prediction: Draw the particles from the proposal density $\boldsymbol{X}_k^i \sim \mathcal{N}(\boldsymbol{X}_k; \hat{\boldsymbol{\mu}}_k^i, \boldsymbol{\Sigma}_k)$ (3.47).
  Particle correction: Update the weights

$$w_k^i \propto w_{k-1}^i \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k^i)\; p(\boldsymbol{X}_k^i|\boldsymbol{X}_{k-1}^i)}{\mathcal{N}(\boldsymbol{X}_k^i; \hat{\boldsymbol{\mu}}_k^i, \boldsymbol{\Sigma}_k)}$$

  if a resampling criterion is satisfied, e.g., $\widehat{\text{ESS}} < N_{\text{th}}$, see (3.19) **then**
    Draw a new set of particles $\{\boldsymbol{X}_k^i\}_{i\in[1,N]}$ using a resampling method, e.g., the multinomial resampling (Algorithm 3.2).
    Reset the weights to $\frac{1}{N}$.
  end if
  Compute the linear state estimate $\hat{\boldsymbol{X}}_k$ (3.16) and its covariance $\boldsymbol{P}_k$ (3.17).
end for
Return the state estimate $\hat{\boldsymbol{X}}_k$ and its covariance $\boldsymbol{P}_k$, $\forall k$.

The proposal density of the WEnKF comes from EnKF equations. For a better readability, the EnKF correction equation is given back:

$$\boldsymbol{X}_k^i = \boldsymbol{X}_{k|k-1}^i + \boldsymbol{K}_k(\boldsymbol{Y}_k - \boldsymbol{H}_k\; \boldsymbol{X}_{k|k-1}^i + \boldsymbol{V}_k^i) \tag{3.43}$$

By replacing the predicted samples $\boldsymbol{X}_{k|k-1}^i$ by its expression $b_k(\boldsymbol{X}_{k-1}^i) + \boldsymbol{W}_k^i$ in the correction equation (3.43), we have:

$$\boldsymbol{X}_k^i = (\boldsymbol{I}_d - \boldsymbol{K}_k\boldsymbol{H}_k)\, b_k(\boldsymbol{X}_{k-1}^i) + \boldsymbol{K}_k\, \boldsymbol{Y}_k + \boldsymbol{K}_k\, \boldsymbol{V}_k^i + (\boldsymbol{I}_d - \boldsymbol{K}_k\boldsymbol{H}_k)\, \boldsymbol{W}_k^i \tag{3.44a}$$

$$:= \hat{\boldsymbol{\mu}}_k^i + \boldsymbol{\gamma}_k^i \tag{3.44b}$$

where $\hat{\boldsymbol{\mu}}_k^i = (\boldsymbol{I}_d - \boldsymbol{K}_k\boldsymbol{H}_k)\, b_k(\boldsymbol{X}_{k-1}^i) + \boldsymbol{K}_k\,\boldsymbol{Y}_k$ is deterministic and $\boldsymbol{\gamma}_k^i = \boldsymbol{K}_k\,\boldsymbol{V}_k^i + (\boldsymbol{I}_d - \boldsymbol{K}_k\boldsymbol{H}_k)\,\boldsymbol{W}_k^i$ regroups the stochastic terms. As the process noise $\boldsymbol{W}_k$ and the perturbation $\boldsymbol{V}_k$ are two centered Gaussian, the random variable $\boldsymbol{\gamma}_k^i$ follows a Gaussian:

$$\boldsymbol{\gamma}_k^i \sim \mathcal{N}\left(\boldsymbol{\gamma}_k;\, 0, \boldsymbol{\Sigma}_k\right) \tag{3.45}$$

where the covariance $\boldsymbol{\Sigma}_k$ has the following form:

$$\boldsymbol{\Sigma}_k = (\boldsymbol{I}_d - \boldsymbol{K}_k\boldsymbol{H}_k)\,\boldsymbol{Q}_k\,(\boldsymbol{I}_d - \boldsymbol{K}_k\boldsymbol{H}_k)^\top + \boldsymbol{K}_k\boldsymbol{R}_k\boldsymbol{K}_k \tag{3.46}$$

The proposal density of the WEnKF is given by:

$$q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}^i, \boldsymbol{Y}_k) = \mathcal{N}(\boldsymbol{X}_k;\, \hat{\boldsymbol{\mu}}_k^i, \boldsymbol{\Sigma}_k) \tag{3.47}$$

The rewriting of the EnKF correction equation (3.44b) allows us to include the WEnKF within the scope of PFs.

3.4 LIKELIHOOD-FREE METHODS: APPROXIMATE BAYESIAN COMPUTATION

Approximate Bayesian Computation (ABC) method is an approach to address inference problems where the likelihood function is unknown, or expensive to calculate. The likelihood may be unavailable for mathematical reasons, e.g., it is not available in an explicit form that would link the state to the observations, or when the observation model is a rather rough approximation of the true model. There are also computational reasons, e.g., the likelihood can be too expensive to calculate. Moreover, ABC methods are used when the normalising constant of the likelihood is unknown or when the likelihood function is partially known [55]. Unfortunately, if the likelihood is unknown, it is difficult to use standard computational tools to sample the conditional density as the optimal filter Bayes equation (2.10) cannot be performed.

ABC method bypasses the computation of the likelihood via the sampling of pseudo-measurements [35, 55]. Assume that there exists a form of measurement model that allows to simulate pseudo-measurements $\boldsymbol{U}_k$ by plugging the state $\boldsymbol{X}_k$ into it. This model may be a probability density, a differential equation, a stochastic process or a noise term-free equation.

The conditional density is approximated by:

$$p(\boldsymbol{X}_k|\mathcal{Y}_k) \approx p^\varepsilon(\boldsymbol{X}_k|\mathcal{Y}_k) = \int p^\varepsilon(\boldsymbol{X}_k, \boldsymbol{U}_k|\mathcal{Y}_k)\; d\boldsymbol{U}_k \tag{3.48}$$

In its most common form, ABC draws inference from the following modified conditional density:

$$p^\varepsilon(\boldsymbol{X}_k, \boldsymbol{U}_k|\mathcal{Y}_k) = \frac{p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\; p(\boldsymbol{U}_k|\boldsymbol{X}_k)\; \mathbb{1}_{\boldsymbol{A}_{\varepsilon,\boldsymbol{Y}_k}}(\boldsymbol{U}_k)}{\int_{\boldsymbol{A}_{\varepsilon,\boldsymbol{Y}_k}\times d}\; p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\; p(\boldsymbol{U}_k|\boldsymbol{X}_k)\; d\boldsymbol{U}_k d\boldsymbol{X}_k} \tag{3.49}$$

where $\mathbb{1}_{\boldsymbol{B}}(.)$ is the indicator function of a given set $\boldsymbol{B}$, ε the tolerance level and $\boldsymbol{U}_k \in \mathbb{R}^{d_y}$ the pseudo-observations. The set $\boldsymbol{A}_{\varepsilon,\boldsymbol{Y}_k}$ corresponds to the set of pseudo-observations which are close in some sense to the true observations $\boldsymbol{Y}_k$. It is formally defined as:

$$\boldsymbol{A}_{\varepsilon,\boldsymbol{Y}_k} = \left\{ \boldsymbol{U} \in \mathbb{R}^{d_Y} \mid \rho\left(s(\boldsymbol{U}_k), s(\boldsymbol{Y}_k)\right) < \varepsilon \right\} \tag{3.50}$$

where $s : \mathbb{R}^{d_y} \to \mathcal{S}$ represents some summary statistics and $\rho : \mathcal{S} \times \mathcal{S} \to \mathbb{R}^+$ a distance function. The idea behind ABC methods is to produce a good approximation of the conditional density by using a representation summary statistic coupled with a small tolerance [55]. The distance function ρ is usually the Euclidean norm $\|.\|_2$ or the Manhattan norm $\|.\|_1$. In the following, the summary statistic s is the identity function. As, in the context of filtering, the observations is only a $\mathbb{R}^{d_y}$-valued data point at each time-step, it is not crucial to summarise the data with a low dimensional summary statistic. If the Euclidean distance is selected, the pseudo-observations set (3.50) is thus simplified as:

$$\boldsymbol{A}_{\varepsilon,\boldsymbol{Y}_k} = \left\{ \boldsymbol{U} \in \mathbb{R}^{d_y} \mid \|\boldsymbol{U}_k - \boldsymbol{Y}_k\|_2 < \varepsilon \right\} \tag{3.51}$$

The tolerance parameter ε must be sequentially adapted, otherwise the filter may abruptly fail if the true observation is an outlier, making the pseudo-observations set $\boldsymbol{A}_{\varepsilon,\boldsymbol{Y}_k}$ empty.

The indicator function $\mathbb{1}_{\boldsymbol{A}_{\varepsilon,\boldsymbol{Y}_k}}(\boldsymbol{U}_k)$ in the approximate conditional density equation (3.49) can be substitute by a smooth kernel function [16]:

$$\mathcal{K}_\varepsilon\left(\boldsymbol{U}_k - \boldsymbol{Y}_k\right) \tag{3.52}$$

where $\mathcal{K}$ is a positive and symmetric kernel (3.22). The indicator function is in fact a uniform kernel. The conditional density $p(\mathcal{X}_k|\mathcal{Y}_k)$ can thus be approximated by:

$$p(\mathcal{X}_k|\mathcal{Y}_k) = p(\boldsymbol{X}_0) \int p(\mathcal{X}_k, \mathcal{U}_k|\mathcal{Y}_k) d\mathcal{U}_k \tag{3.53}$$

$$\approx p(\boldsymbol{X}_0) \prod_{k\geq 1} \left[\int \mathcal{K}_\varepsilon(\boldsymbol{Y}_k, \boldsymbol{U}_k) p(\boldsymbol{U}_k|\boldsymbol{X}_k) d\boldsymbol{U}_k \right] p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}) \tag{3.54}$$

In practice, heavy-tailed kernels are used as e.g., the Gaussian and Cauchy kernels [16].

ABC method was first introduced as a rejection technique [73, 82]. The resulting ABC rejection sampler is summarised in Algorithm 3.7. A Markov chain Monte Carlo (MCMC)-ABC algorithm was also introduced [56]. Later, an ABC scheme for filtering was proposed [43]. The ABC filter in its bootstrap form is summarised in Algorithm 3.8.

The tolerance parameter ε plays a significant role in the convergence of ABC method [43]. The choice of this parameter depends on the considered ABC algorithm. Several choices for the tolerance parameter were proposed [16, 22, 43]. For example to avoid degeneracy, the parameter can be adaptively selected according to a criterion based on the approximate effective sample size (3.19) (e.g., see [22, 43]) or according to the MISE [16].

Algorithm 3.7 ABC rejection sampler

Given the state $\boldsymbol{X}_{k-1}$,
for each $i \in [1, N]$ **do**
  Repeat
  Generate the state $\boldsymbol{X}'_k$ from the prior density $p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1})$
  Generate the pseudo-measurement $\boldsymbol{U}_k$ from the density $p(\boldsymbol{U}_k|\boldsymbol{X}_k)$.
  until $\rho\left(s(\boldsymbol{U}_k), s(\boldsymbol{Y}_k)\right) \leq \varepsilon$
  Set $\boldsymbol{X}^i_k = \boldsymbol{X}'_k$
end for
Return the state samples $\boldsymbol{X}^i_k$, $\forall i$.

Algorithm 3.8 ABC filter

Initialisation: The initial particle set is drawn as $\{\boldsymbol{X}^i_0\}_{i\in[1,N]}$ using the initial density $p(\boldsymbol{X}_0)$ and the initial weights set $\{w^i_0\}_{i\in[1,N]}$ is taken equal to $\frac{1}{N}$.
for each time-step k **do**
  Prediction: Draw the particles from the transition density $\boldsymbol{X}^i_k \sim p(\boldsymbol{X}_k|\boldsymbol{X}^i_{k-1})$.
  For each particle $\boldsymbol{X}^i_k$, simulate a pseudo-measurement $\boldsymbol{U}^i_k$ using the density $p(\boldsymbol{U}_k|\boldsymbol{X}_k)$.
  Correction: Update the weights $w^i_k \propto w^i_{k-1}\, \mathcal{K}_\varepsilon\left(\boldsymbol{Y}_k - \boldsymbol{U}^i_k\right)$.
  Compute the state estimate $\hat{\boldsymbol{X}}_k$ (3.16) and its covariance $\boldsymbol{P}_k$ (3.17).
  if a resampling criterion is satisfied, e.g., $\widehat{ESS} < N_{\text{th}}$, see (3.19) **then**
    Draw a new set of particles $\{\boldsymbol{X}^i_k\}_{i\in[1,N]}$ using a resampling method, e.g., the multinomial resampling (Algorithm 3.2).
    Reset the weights to $\frac{1}{N}$.
  end if
end for
Return the state estimate $\hat{\boldsymbol{X}}_k$ and its covariance $\boldsymbol{P}_k$, $\forall k$.

3.5 MONTE CARLO ERROR

In Section 3.3.1, several sources of Monte Carlo errors were introduced. As discussed in the paragraph on the degeneracy problem, the variance of the weights increases over time. To address the degeneracy problem, two solutions were proposed: the resampling and/or the choice of the proposal density. However, resampling introduces dependency between particles and causes sampling impoverishment phenomenon, and the choice of the proposal density is not straightforward. Thus, the resampling is performed only when the approximate effective sample size $\widehat{\mathrm{ESS}}$ is under a given threshold N_{th}, and the proposal density is often taken equal to the prior density (bootstrap filter). For bootstrap filters, Monte Carlo errors arise when the support of the proposal density does not contain the support of the prior density, as shown in Figure 3.2b. Thus, the performance of Monte Carlo approaches depends on the choice of the proposal density and on the consistency between the likelihood and the prior density.

The consistency is formally defined by the following integral:

$$\int p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k \tag{3.55}$$

Two criteria are derived from the study of the consistency: the asymptotic variance of the unnormalised weights and the discrepancy of estimation.

3.5.1 *Asymptotic Variance of the Unnormalised Weights*

Proposition 1. *The asymptotic variance of the unnormalised weights is given by:*

$$\mathbb{V} = \frac{\int \tilde{w}^2\, q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)\, d\boldsymbol{X}_k}{\left(\int \tilde{w}\, q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)\, d\boldsymbol{X}_k\right)^2} - 1 \tag{3.56}$$

where $\tilde{w} = \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)}$ *are the unnormalised weights.*

Proof. The variance of the unnormalised weights is:

$$\mathbb{V}\mathrm{ar}_q(\tilde{w}) = \int \tilde{w}^2\, q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)\, d\boldsymbol{X}_k - \left(\int \tilde{w}\, q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)\, d\boldsymbol{X}_k\right)^2 \tag{3.57a}$$

$$= \left(\int \tilde{w}\; q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)\; d\boldsymbol{X}_k\right)^2 \left(\frac{\int \tilde{w}^2\, q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)\, d\boldsymbol{X}_k}{\left(\int \tilde{w}\, q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)\, d\boldsymbol{X}_k\right)^2} - 1\right) \tag{3.57b}$$

where $\tilde{w} = \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1},\boldsymbol{Y}_k)}$ are the unnormalised weights. Then,

$$\frac{\mathbb{V}\mathrm{ar}_q(\tilde{w})}{\left(\int \tilde{w}\; q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1},\boldsymbol{Y}_k)\; d\boldsymbol{X}_k\right)^2} = \left(\frac{\int \tilde{w}^2\, q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1},\boldsymbol{Y}_k)\, d\boldsymbol{X}_k}{\left(\int \tilde{w}\, q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1},\boldsymbol{Y}_k)\, d\boldsymbol{X}_k\right)^2} - 1\right) \tag{3.58}$$

By taking,

$$\mathbb{E}_q(\tilde{w}) = \int \tilde{w}\; q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1},\boldsymbol{Y}_k)\; d\boldsymbol{X}_k \tag{3.59}$$

the asymptotic variance of the unnormalised weights $\mathbb{V}$ comes:

$$\mathbb{V} = \frac{\mathbb{V}\mathrm{ar}_q(\tilde{w})}{(\mathbb{E}_q(\tilde{w}))^2} \tag{3.60}$$

This variance is a squared coefficient of variance (or a normalised variance). □

By replacing the unnormalised weights by their expression in the asymptotic variance, it comes:

$$\mathbb{V} = \frac{\int \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k)^2\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})^2}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1},\boldsymbol{Y}_k)}\, d\boldsymbol{X}_k}{\left(\int p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k\right)^2} - 1 \tag{3.61}$$

The asymptotic variance $\mathbb{V}$ is equal to zero if the support of the proposal density recovers the one of the conditional density, i.e., $q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1},\boldsymbol{Y}_k) \propto p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\; p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})$. As discussed previously, the choice of the proposal density is crucial to design robust particle filters.

When the proposal density is taken equal to the prior density, the asymptotic variance becomes:

$$\mathbb{V} = \frac{\int \tilde{w}^2 p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k}{\left(\int \tilde{w} p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k\right)^2} - 1 \tag{3.62}$$

where the unnormalised weights are equal to $\tilde{w} = p(\boldsymbol{Y}_k|\boldsymbol{X}_k)$.

The asymptotic variance of the unnormalised weights is small when:

- the support of the prior density $p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})$ includes the support of the likelihood $p(\boldsymbol{Y}_k|\boldsymbol{X}_k)$
- the support of the proposal density recovers the support of the conditional density, i.e., $q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k) \propto p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\ p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})$

Guaranteeing a finite and small asymptotic variance of the unnormalised weights ensures that the Monte Carlo estimate is accurate, which reduces the degeneracy problem.

3.5.2 *Discrepancy*

The inconsistency between two densities can be quantified by the discrepancy criterion $\gamma_k \geq 1$. The discrepancy characterizes the consistency between the current measurements and the predicted measurements. The discrepancy is large when the variance of the observation noise is low, and when the likelihood is inconsistent with the prior density. The criterion is inversely proportional to the consistency between the likelihood and the prior density. It is formally defined as:

$$\gamma_k \triangleq \frac{\sup\limits_{\boldsymbol{X}_k \in \mathbb{R}^d} \tilde{w}}{\int \tilde{w} q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)\, d\boldsymbol{X}_k} \tag{3.63}$$

By replacing the unnormalised weights in the discrepancy (3.63), it comes:

$$\gamma_k = \frac{\sup\limits_{\boldsymbol{X}_k \in \mathbb{R}^d} \left(\frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\ p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \right)}{\int \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\ p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)\, d\boldsymbol{X}_k} \tag{3.64a}$$

$$= \frac{\sup\limits_{\boldsymbol{X}_k \in \mathbb{R}^d} \left(p(\boldsymbol{Y}_k|\boldsymbol{X}_k) \frac{p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \right)}{\int p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\ p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k} \tag{3.64b}$$

When the proposal density $q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)$ is equal to the prior density $p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})$, the discrepancy criterion can be simplified by:

$$\gamma_k = \frac{\sup\limits_{\boldsymbol{X}_k \in \mathbb{R}^d} p(\boldsymbol{Y}_k|\boldsymbol{X}_k)}{\int p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\ p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k} \tag{3.65}$$

In this case, the discrepancy corresponds to the inverse of the probability of acceptance in acceptance-rejection sampling algorithms [26]:

$$\gamma_k = \frac{1}{\mathbb{P}_A} \tag{3.66}$$

where

$$\mathbb{P}_A = \frac{\int p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\ p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k}{\sup\limits_{\boldsymbol{X}_k \in \mathbb{R}^d} p(\boldsymbol{Y}_k|\boldsymbol{X}_k)} \tag{3.67}$$

If the probability of acceptance is close to 1, the discrepancy is small, and therefore a lot of particles have a non-zero weight. On the contrary, if the probability $\mathbb{P}_A$ is small, the discrepancy is large, which indicates that the support of the likelihood does not recover the support of the prior density.

3.6 SUMMARY

To cope with non-Gaussian and non-linear cases, Monte Carlo approaches were introduced in Section 3.3.1. These methods empirically describe the non-linear propagation and update of the conditional density. More advanced and robust particle filters were introduced based on kernel regularisation (Section 3.3.2), on conditionally linear models (Section 3.3.3), or on EnKF equations (Section 3.3.4). A likelihood-free method called ABC was introduced in Section 3.4. Although it is robust to non-linearities and ambiguities, these filters may diverge either because of a suboptimal choice of proposal density, or because of the inconsistency between the likelihood and the prior density.

4

UNDERWATER NAVIGATION

This chapter provides elements on the work's methodology, in particular concerning the simulations and their interpretations. Simulations are presented in order to illustrate the theoretical results and to evaluate the practical impact of the work.

Section 4.1 introduces the inertial navigation equations used to evaluate the state estimation algorithms (Chapters 5 and 6). The inertial navigation equation described the kinematic parameters of a vehicle over time. Time-varying linear models associated with negligible process noise were chosen in order to focus on the impact of measurements ambiguities and non-linearities. The considered models are the inertial measurement unit error drift and the double integrator.

Section 4.2 introduces the multi-beam telemeter and atomic gravimeter, and the associated observation models are described. The work focuses on Terrain Aided Navigation (TAN) application. TAN constitutes a severely ambiguous problem and allows estimation algorithms to be compared in a significant context. Furthermore, it belongs to a larger class of problems that consist of state estimation from vector field measurements. In the context of underwater TAN, two sensors provided the measurements: the multi-beam telemeter (see Section 4.2.1) and the atomic gravimeter (see Section 4.2.2). The sensor fusion strategy is explained in Section 4.2.3.

Section 4.3 presents TAN scenarios. The bathymetric and gravimetric maps are introduced.

Section 4.4 defines numerical criteria used to interpret the simulations: the Root Mean Square Error (RMSE), the Posterior Cramér-Rao Bound (PCRB) and the non-convergence rate.

4.1 DYNAMICAL MODELS FOR STATE ESTIMATION

Consider a vehicle whose state vector consists of its geographical position, velocity, and attitude in the Earth frame:

$$\boldsymbol{X} = \left[\boldsymbol{p}^\top, \boldsymbol{V}^\top, \boldsymbol{\Psi}^\top\right]^\top \in \mathbb{R}^9 \tag{4.1}$$

where $\boldsymbol{p} = [p_\lambda, p_\phi, p_h]^\top$ is the geographical position (respectively latitude (rad), longitude (rad), altitude (m)), $\boldsymbol{V} = [V_N, V_E, V_D]^\top \in \mathbb{R}^3$ is the velocity vector ($\mathrm{m\,s^{-1}}$), and $\boldsymbol{\Psi} = [\psi, \theta, \varphi]^\top$ is the attitude (euler angles in rad, respectively heading, pitch, roll). These values

are unknown and are usually estimated by an Inertial Measurement Unit (IMU). The conventional reference system is summarised in Figure 4.1. The inertial frame is fixed and Earth-centered. The Earth frame is deduced from the inertial frame through the rotation of the Earth. The navigation frame moves across the Earth's surface at the same time as the vehicle.

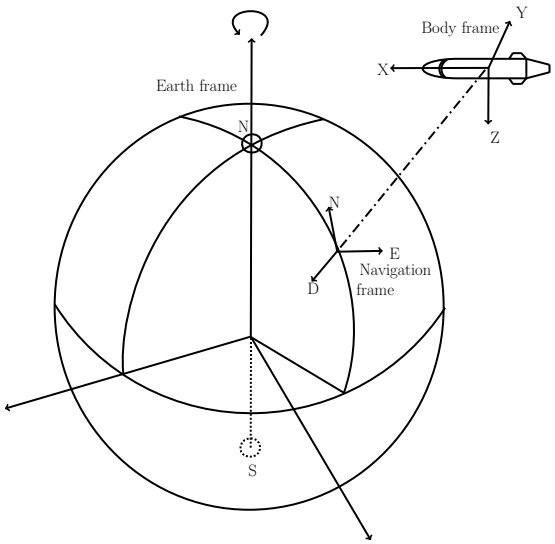

Figure 4.1: Main Earth reference frames.

4.1.1 *Inertial Navigation Equations in the Navigation Frame*

The IMU measures the vehicle's proper acceleration $\boldsymbol{Y}^a$ and angular rate $\boldsymbol{Y}^\omega$ in the inertial frame, via respectively three accelerometers and three gyrometers (directed according to the directions of the navigation frame, see Figure 4.1). These measurements are integrated to compute the accelerations $\dot{\boldsymbol{V}}$ and angular rates $\left[\dot{\boldsymbol{R}}_{n2b}\right]$. The position is then obtained from the double integration of the acceleration.

The acceleration in the Earth frame is given by the combination of the vehicle's proper acceleration $\boldsymbol{Y}^a$ and a gravity model. The following simplified gravity model in the World Geodetic System 1984 (WGS 84) is used:

$$g = 9.7803 + 0.0519 \sin^2(p_\lambda) - 3.08 \times 10^{-6}\, p_h \tag{4.2}$$

where g is expressed in $\mathrm{m\,s^{-2}}$. The acceleration is modelled by the following equation:

$$\dot{\boldsymbol{V}} = [\boldsymbol{R}_{n2b}]\ \boldsymbol{Y}^a + g - (2\,\boldsymbol{\omega}_{ie} + \boldsymbol{\rho}) \wedge \boldsymbol{V} \tag{4.3}$$

where

- $[\boldsymbol{R}_{n2b}]$ is the attitude matrix. The matrix allows switching from the navigation frame to the body frame and is equal to:

$$[\boldsymbol{R}_{n2b}] = \begin{pmatrix} \cos\theta\cos\psi & \cos\theta\sin\psi & -\sin\theta \\ -\cos\varphi\sin\psi + \sin\varphi\sin\theta\cos\psi & \cos\varphi\cos\psi + \sin\varphi\sin\theta\sin\psi & \sin\varphi\cos\theta \\ \sin\varphi\sin\psi + \cos\varphi\sin\theta\cos\psi & -\sin\varphi\cos\psi + \cos\varphi\sin\theta\sin\psi & \cos\varphi\cos\theta \end{pmatrix} \tag{4.4}$$

  The attitude matrix verifies the following property: $[\boldsymbol{R}_{n2b}] = [\boldsymbol{R}_{b2n}]^\top$.

- $\boldsymbol{\omega}_{ie} + \boldsymbol{\rho}$ is the Earth rotation rate. $\boldsymbol{\rho}$ is the angular velocity of the navigation frame with respect to the Earth frame and $\boldsymbol{\omega}_{ie}$ is the velocity of the Earth rotation on itself. These terms are defined in the navigation frame by:

$$\boldsymbol{\rho} = \begin{bmatrix} \dot{p}_\phi \cos p_\lambda \\ -\dot{p}_\lambda \\ \dot{p}_\phi \sin p_\lambda \end{bmatrix} \tag{4.5}$$

$$\boldsymbol{\omega}_{ie} = \begin{bmatrix} \omega_0 \cos p_\lambda \\ 0 \\ -\omega_0 \sin p_\lambda \end{bmatrix} \tag{4.6}$$

  where $\omega_0 = 7.29 \times 10^{-5}$ rad s^{-1}.

- The notation $\wedge$ denotes the cross product.

The angular rate $\boldsymbol{Y}^\omega$ is compensated by the Earth rotation rate, which leads to the angular rate $\left[\dot{\boldsymbol{R}}_{n2b}\right]$ modelled by:

$$\left[\dot{\boldsymbol{R}}_{n2b}\right] = -[\boldsymbol{Y}^\omega\times]\,[\boldsymbol{R}_{n2b}] + [\boldsymbol{R}_{n2b}]\,[(\boldsymbol{\omega}_{ie} + \boldsymbol{\rho})\times] \tag{4.7}$$

where $[\boldsymbol{Y}^\omega\times] = \begin{pmatrix} 0 & -r & q \\ r & 0 & -p \\ -q & p & 0 \end{pmatrix}$ is the anti-symmetric matrix of the vector $\boldsymbol{Y}^\omega = [p, q, r]^\top$.

The IMU is coupled with a double integrator which integrates the accelerations (4.3) and angular rates (4.7) to provide an estimate of the vehicle's state, denoted $\tilde{\boldsymbol{X}}^{\text{IMU}}$.

IMU accelerometers and gyrometers measurements are uncertain and can be modelled by the following equations:

$$\boldsymbol{Y}^a = \boldsymbol{b}^a + (1 + \boldsymbol{K}^a)\,\boldsymbol{a} + \boldsymbol{v}^a \tag{4.8a}$$

$$\boldsymbol{Y}^\omega = \boldsymbol{b}^\omega + (1 + \boldsymbol{K}^\omega)\,\boldsymbol{\omega} + \boldsymbol{v}^\omega \tag{4.8b}$$

where $\boldsymbol{a} \in \mathbb{R}^3$ and $\boldsymbol{\omega} \in \mathbb{R}^3$ are the actual vehicle proper acceleration and angular rate in the inertial frame. $\boldsymbol{K}^a \in \mathbb{R}^3$ and $\boldsymbol{K}^\omega \in \mathbb{R}^3$ are scale factors. Noises $\boldsymbol{v}^a \in \mathbb{R}^3$ and $\boldsymbol{v}^\omega \in \mathbb{R}^3$ are Wiener processes. Vectors $\boldsymbol{b}^a$ and $\boldsymbol{b}^\omega$ are respectively the accelerometers and gyrometers bias. They are modelled with a first order Markov chain:

$$\dot{\boldsymbol{b}}^a = -\frac{1}{\tau_a}\boldsymbol{b}^a + \boldsymbol{\nu}^a \tag{4.9}$$

$$\dot{\boldsymbol{b}}^\omega = -\frac{1}{\tau_\omega}\boldsymbol{b}^\omega + \boldsymbol{\nu}^\omega \tag{4.10}$$

where τ_a and τ_ω are the correlation period of $\boldsymbol{b}^a$ and $\boldsymbol{b}^\omega$ respectively, and $\boldsymbol{\nu}^a$ and $\boldsymbol{\nu}^\omega$ are Wiener processes. Refined models include sensors misalignments, or scale factors in the IMU measurements equations (4.8). By iterative integration of IMU measurements, the IMU state drifts from the actual state and may rapidly provide poor navigation information.

4.1.2 *Inertial Errors Model*

The IMU errors are the difference between the actual state and the IMU state. They can be modelled as a state vector $\delta\boldsymbol{X}$ which includes an estimation of the accelerometers and gyrometers bias:

$$\delta\boldsymbol{X} = \left[\delta\boldsymbol{x}^\top, \delta\boldsymbol{V}^\top, \delta\boldsymbol{\Psi}^\top, (\boldsymbol{b}^a)^\top, (\boldsymbol{b}^\omega)^\top\right]^\top \in \mathbb{R}^{15} \tag{4.11}$$

where $\delta\boldsymbol{x}$ is the metric position error, $\delta\boldsymbol{V}$ the velocity error, $\delta\boldsymbol{\Psi}$ the attitude angles errors. The scale factors $\boldsymbol{K}^\omega$ and $\boldsymbol{K}^a$ are not estimated.

The relations between the metric position error and the position error are given by:

$$\begin{cases} \delta x_N = (R_\lambda + p_h)\,\delta p_\lambda \\ \delta x_E = (R_\phi + p_h)\,cos(p_\lambda)\,\delta p_\phi \\ \delta x_D = -\delta p_h \end{cases} \tag{4.12}$$

where R_λ is the radius of curvature of the Earth in the meridian plane and R_ϕ is the large normal of the ellipsoid.

Inertial errors equations are obtained by first-order differentiation of the inertial model (equations (4.3) and (4.7)) with respect to the coordinates of the problem. The inertial errors model can be written as follows:

$$\begin{cases} \delta\dot{\boldsymbol{\Psi}} &= -(\boldsymbol{\rho} + \boldsymbol{\omega}_{ie}) \wedge \delta\boldsymbol{\Psi} - [\boldsymbol{R}_{b2n}]\,(\boldsymbol{b}^\omega + \boldsymbol{v}^\omega) \\ \delta\dot{\boldsymbol{V}} &= -\delta\boldsymbol{\Psi} \wedge \boldsymbol{f} + [\boldsymbol{R}_{b2n}]\,(\boldsymbol{b}^a + \boldsymbol{v}^a) + \delta\boldsymbol{g} - (\boldsymbol{\rho} + 2\,\boldsymbol{\omega}_{ie}) \wedge \delta\boldsymbol{V} \\ \delta\dot{\boldsymbol{x}} &= \delta\boldsymbol{V} - \boldsymbol{\rho} \wedge \delta\boldsymbol{x} \end{cases} \tag{4.13}$$

where

- $\delta\boldsymbol{\Psi} = [\delta\psi, \delta\theta, \delta\varphi]^\top$ is the vector of the attitude angles errors

- $\boldsymbol{\rho} = [\rho_N, \rho_E, \rho_D]^\top$ is the angular velocity of the Earth rotation expressed in the navigation frame.
- $\boldsymbol{f} = [f_N, f_E, f_D]^\top$ is the vector of proper accelerations projected in the navigation frame.

The inertial errors dynamics can be described by a linearised model:

$$\delta \boldsymbol{X}_k = \boldsymbol{A}_k \, \delta \boldsymbol{X}_{k-1} + \boldsymbol{B}_k \, \boldsymbol{W}_k \tag{4.14}$$

where matrices $\boldsymbol{A}_k \in \mathbb{R}^{15\times15}$ and $\boldsymbol{B}_k \in \mathbb{R}^{15\times12}$ depend on the state, the gravity model, the measured proper acceleration $\boldsymbol{Y}_k^a$, and the IMU characteristics. These matrices are given by:

$$\boldsymbol{A}_k = \begin{bmatrix} \mathbf{1}_3 + \boldsymbol{F}_{XX}\Delta t & \mathbf{1}_3\Delta t & 0 & 0 & 0 \\ \boldsymbol{F}_{VX}\Delta t & \mathbf{1}_3 + \boldsymbol{F}_{VV}\Delta t & \boldsymbol{F}_{V\Psi}\Delta t & [\boldsymbol{R}_{b2n}]\,\Delta t & 0 \\ 0 & 0 & \mathbf{1}_3 + \boldsymbol{F}_{\Psi\Psi}\Delta t & 0 & -[\boldsymbol{R}_{b2n}]\,\Delta t \\ 0 & 0 & 0 & \mathbf{1}_3\left(1 - \frac{1}{\tau_a}\Delta t\right) & 0 \\ 0 & 0 & 0 & 0 & \mathbf{1}_3\left(1 - \frac{1}{\tau_\omega}\Delta t\right) \end{bmatrix} \tag{4.15a}$$

$$\boldsymbol{B}_k = \begin{bmatrix} 0 & 0 & 0 & 0 \\ [\boldsymbol{R}_{b2n}]\sqrt{\Delta t} & 0 & 0 & 0 \\ 0 & -[\boldsymbol{R}_{b2n}]\sqrt{\Delta t} & 0 & 0 \\ 0 & 0 & \mathbf{1}_3\sqrt{\Delta t} & 0 \\ 0 & 0 & 0 & \mathbf{1}_3\sqrt{\Delta t} \end{bmatrix} \tag{4.15b}$$

where Δt is the time-step, and $\boldsymbol{F}_{XX}$, $\boldsymbol{F}_{VX}$, $\boldsymbol{F}_{VV}$, $\boldsymbol{F}_{V\Psi}$ and $\boldsymbol{F}_{\Psi\Psi}$ are defined by:

$$\boldsymbol{F}_{XX} = \begin{bmatrix} 0 & \rho_D & -\rho_E \\ -\rho_D & 0 & \rho_N \\ \rho_E & -\rho_N & 0 \end{bmatrix} \tag{4.16}$$

$$\boldsymbol{F}_{VX} = \begin{bmatrix} -\left(\omega_s^N\right)^2 & 0 & 0 \\ 0 & -\left(\omega_s^E\right)^2 & 0 \\ 0 & 0 & 2\left(\omega_s^D\right)^2 \end{bmatrix} \tag{4.17}$$

$$\boldsymbol{F}_{VV} = \begin{bmatrix} 0 & \left(\rho_D + 2\,\omega_{ie}^D\right) & -\rho_E \\ -\left(\rho_D + 2\,\omega_{ie}^D\right) & 0 & \left(\rho_N + 2\,\omega_{ie}^N\right) \\ \rho_E & -\left(\rho_N + 2\,\omega_{ie}^N\right) & 0 \end{bmatrix} \tag{4.18}$$

$$\boldsymbol{F}_{V\Psi} = \begin{bmatrix} 0 & -f_D & f_E \\ f_D & 0 & -f_N \\ -f_E & f_N & 0 \end{bmatrix} \tag{4.19}$$

$$\boldsymbol{F}_{\Psi\Psi} = \begin{bmatrix} 0 & \left(\rho_D + \omega_{ie}^D\right) & -\rho_E \\ -\left(\rho_D + \omega_{ie}^D\right) & 0 & \left(\rho_N + \omega_{ie}^N\right) \\ \rho_E & -\left(\rho_N + \omega_{ie}^N\right) & 0 \end{bmatrix} \tag{4.20}$$

where $\left[\omega_S^N, \omega_S^E, \omega_S^D\right] = \left[\sqrt{\frac{g}{R_\lambda}}, \sqrt{\frac{g}{R_\phi}}, \sqrt{\frac{g}{\sqrt{R_\phi R_\lambda}}}\right]$ is the vector of the Schuler frequency on each axis. R_λ is the radius of curvature and R_ϕ is the semi-major axis of the ellipsoid in the WGS 84 model.

The process noise $\boldsymbol{W}_k \in \mathbb{R}^{12}$ is a vector composed of Gaussian white noises associated with v^a and v^ω (see equation (4.8)), and ν^a and ν^ω (see equations (4.9) and (4.10)).

Note that the model (4.14) is independent from the vehicle dynamical model, since the dynamics are directly measured and integrated by the IMU and integrator system. A complete derivation of inertial equations and inertial errors model can be found in Britting [14] and Dahia [19].

4.1.3 *Inertial Measurement Unit Hybridisation Scheme*

The inertial state has to be corrected using additional measurements (e.g., provided by GNSS, radio navigation, or optical sensors). Knowledge of the IMU state drift allows the problem to be formalised as a state estimation scheme. The use of an estimator makes it possible to estimate the IMU errors and to retrieve the actual state as a corrected IMU state. This estimation process is called IMU hybridisation. The estimator is referred to as hybridisation filter and to the additional sensors as hybridisation sensors. Figure 4.2 illustrates the IMU hybridisation scheme.

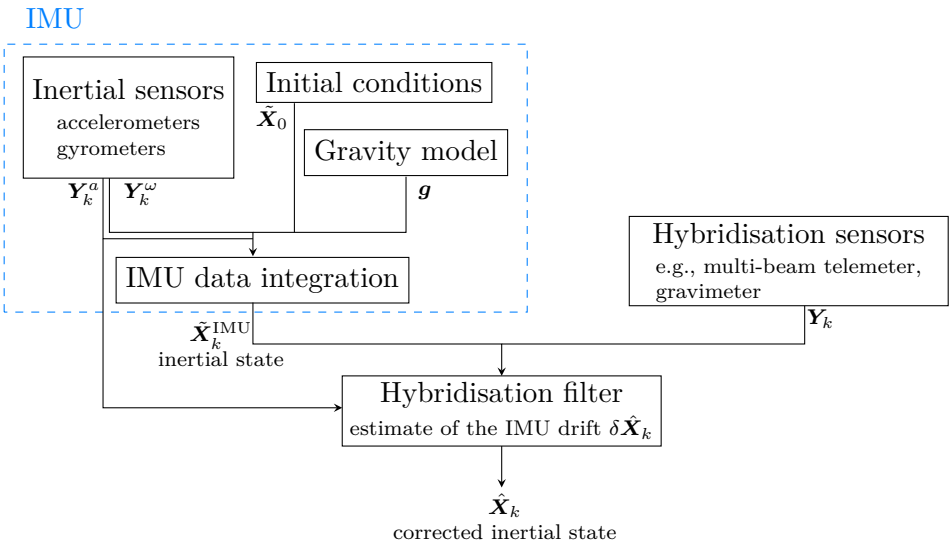

Figure 4.2: IMU hybridisation scheme: the IMU measures the vehicle accelerations and angular rates that are integrated to produce an IMU state estimate. Since this estimate drifts, an hybridisation of the IMU with additional measurements is needed, which provides a corrected state estimate.

4.1.4 *Double Integrator*

The double integrator dynamical model is defined by:

$$\dot{\boldsymbol{X}} = \begin{bmatrix} \mathbf{0}_3 & \mathbf{1}_3 \\ \mathbf{0}_3 & \mathbf{0}_3 \end{bmatrix} \boldsymbol{X} + \boldsymbol{W} \tag{4.21}$$

where the state $\boldsymbol{X} = \left[\boldsymbol{p}^\top, \boldsymbol{V}^\top\right]^\top$ is comprised of three Cartesian positions $\boldsymbol{p} = [p_x, p_y, p_z]$ and three velocities $\boldsymbol{V} = [V_x, V_y, V_z]$, and $\boldsymbol{W} \in \mathbb{R}^6$ is the process noise. In practice, it can be approximated by the Euler method, for a time-step $\Delta t \in \mathbb{R}^+$:

$$\boldsymbol{X}_k = \begin{bmatrix} \mathbf{1}_3 & \Delta t\, \mathbf{1}_3 \\ \mathbf{0}_3 & \mathbf{1}_3 \end{bmatrix} \boldsymbol{X}_{k-1} + \boldsymbol{W}_k \tag{4.22}$$

Equation (4.22) models in a simplified way the position drift of an Autonomous Underwater Vehicle (AUV) IMU that would start with a position and velocity error and a negligible velocity drift. This approximation is meaningful when the IMU gyrometers are sufficiently accurate for short duration scenarios.

4.2 TERRAIN AIDED NAVIGATION

To correct the navigation drift, IMU can be combined with other external sensors. A common aiding source is the Global Positioning System (GPS) but resurfacing for GPS is often excluded for discretion requirements and also because it can be easily jammed. This is especially true for military-grade AUVs. Terrain Aided Navigation (TAN) provides a drift-free navigation tool for underwater operations. This method aims to retrieve the vehicle current state by matching a terrain profile obtained from a sensor with a profile reconstructed from an embedded map of the operation area. The availability of underwater maps as well as the appearance of precise gravimeters make the method opportune for underwater operations.

Usually, the multi-beam telemeter is used for underwater navigation applications [58, 83]. This sensor provides a series of depth measurements along the AUV trajectory. If the terrain contains sufficient information, this sensor can be used to retrieve the state of the AUV using an embedded numerical seabed map. Although multi-beam telemeter aided navigation is able to accurately retrieve the vehicle position, the velocity is often challenging to be precisely estimated. In order to enhance the accuracy of the velocity estimation, we propose to fuse the multi-beam telemeter with an atomic gravimeter. The atomic gravimeter is a promising absolute sensor for underwater navigation. It provides an absolute and accurate gravity measurement. The gravimeter is associated with a gravity anomalies map.

Multi-beam telemeter and gravimeter sensors, as well as the resulting measurements equations, are described in Section 4.2.1 and Section 4.2.2 respectively.

4.2.1 *Multi-beam Telemeter*

The multi-beam telemeter provides a series of depth measurements along the AUV position as shown in Figure 4.3.

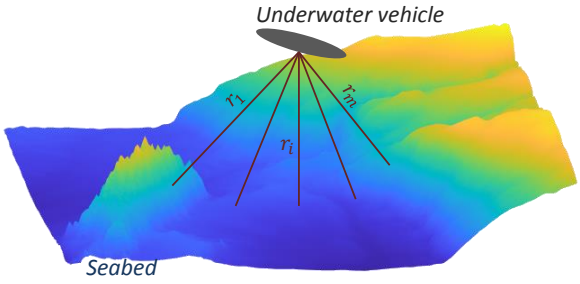

Figure 4.3: Multi-beam telemeter scheme.

At each time-step, the telemeter measurements are comprised of m distances between the AUV and the seabed, which correspond to the number of beams. The measurements vector is given by:

$$\boldsymbol{Y}^b = [r_1, \dots, r_m]^\top \tag{4.23}$$

where r_i records the distance between the AUV and the seabed for the i-th beam.

The observation equation depends of the dynamical model considered.

Multi-beam Telemeter Observation Equation for the Inertial Errors Model

The measurement equation is constructed via a projection in the Cartesian coordinate system (see Meduna [57] for details). For each $i \in [\![1, m]\!]$,

$$r_i = \sqrt{\left(p_\lambda - p_\lambda^i\right)^2 + \left(p_\phi - p_\phi^i\right)^2 + \left(p_h - \mathrm{map}_{\mathrm{mb}}(p_\lambda^i, p_\phi^i)\right)^2} + \nu_b^i \tag{4.24}$$

where the vector $[p_\lambda, p_\phi, p_h]$ is the true geographical position of the AUV. The true position is given by the following relations:

$$\begin{cases} p_\lambda = \tilde{p}_\lambda + \delta p_\lambda \\ p_\phi = \tilde{p}_\phi + \delta p_\phi \\ p_h = \tilde{p}_h + \delta p_h \end{cases} \tag{4.25}$$

where $[\tilde{p}_\lambda, \tilde{p}_\phi, \tilde{p}_h]$ is the position given by the IMU and $[\delta p_\lambda, \delta p_\phi, \delta p_h]$ is the vector of position errors that will be estimated by the filter. Equation (4.24) can be rewritten as follows:

$$r_i = \sqrt{\left(\tilde{p}_\lambda + \delta p_\lambda - p_\lambda^i\right)^2 + \left(\tilde{p}_\phi + \delta p_\phi - p_\phi^i\right)^2 + \left(\tilde{p}_h + \delta p_h - \mathrm{map}_{\mathrm{mb}}(p_\lambda^i, p_\phi^i)\right)^2} + \nu_b^i \tag{4.26}$$

where map_{mb} is the numerical map of the operation area and ν_b^i is a white noise of covariance $\boldsymbol{R}^b$. The range r_i is computed by determining the intersection point of the i^{th} beam direction vector with the terrain $\boldsymbol{p}^i = [p_\lambda^i, p_\phi^i, p_h^i]^\top$ where $p_h^i = \text{map}_{\text{mb}}(p_\lambda^i, p_\phi^i)$ (see Figure 3.4 in [57]). The intersection point is unknown in practice and is analytically linked to the beam r_i through trigonometric relations. In the 2D case, denoting Θ the known angle between the beam and the vertical axis, the coordinates can be written as follows: $[p_\lambda^i, p_\phi^i]^\top = [\tilde{p}_\lambda + \sin\Theta\ r_i\ ,\ \tilde{p}_\phi + \cos\Theta\ r_i]^\top$. Using these relations in the observation equation (4.26) results in an implicit equation where r_i is on both sides of the equal sign.

Multi-beam Telemeter Observation Equation for the Double Integrator Model

Similarly, for each $i \in [\![1, m]\!]$, the beam equation is defined by:

$$r_i = \sqrt{(p_x - p_x^i)^2 + (p_y - p_y^i)^2 + (p_z - \text{map}_{\text{mb}}(p_x^i, p_y^i))^2} + \nu_b^i \tag{4.27}$$

where map_{mb} is the bathymetric map of the operation area (see Figure 4.7) and ν_b^i is a white noise of covariance $\boldsymbol{R}^b$. The position vector $[p_x, p_y, p_z]$ will be estimated by the filter. The bathymetric map is a nonlinear function of $\mathbb{R}^2$ in $\mathbb{R}$, taking as input the position along axis x and y and giving as output the elevation of the terrain. The range r_i is computed by determining the intersection point of the i^{th} beam direction vector with the terrain $\boldsymbol{p}^i = [p_x^i, p_y^i, p_z^i]^\top$ where $p_z^i = \text{map}_{\text{mb}}(p_x^i, p_y^i)$.

As the intersection point is unknown in practice, the multi-beam telemeter observation equation (4.27) is computed through numerical approximations (e.g., via grid search methods or ray tracing code functions). The measurements noise accounts for these approximations in addition to sensor and map errors. However, as the numerical approximations introduce some unknown sampling noise that are difficult to control, the law of the observation noise is complex to infer.

4.2.2 *Gravimeter*

The atomic gravimeter measures the absolute value of the gravity by monitoring the free-fall acceleration of ultra-cold atoms thanks to atom interferometry [13, 18] (see Figure 4.4). The measurements are absolute and accurate, which means that the sensors does not need any calibration and provides a gravity evaluation with a very low level noise (on the order of 10^{-2} mGal, where 1 mGal is equal to 10^{-5} $\text{m}\,\text{s}^{-2}$).

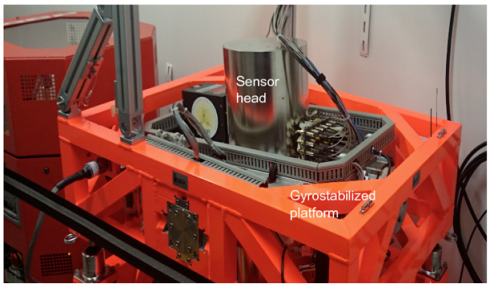

Figure 4.4: Cold atom gravimeter GIRAFE 2 and its gyrostabilizing platform.

Gravimeter Observation Equation for the Inertial Errors Model

The gravimeter observation equation is derived from the following physical equation written in the WGS 84 model:

$$Y^g = g_{nom}(p_\lambda, p_h) + map_{ga}(p_\phi, p_\lambda) + \ddot{p}_h + Eötvös(V_E, p_\lambda) + \nu \tag{4.28}$$

where $\boldsymbol{p} = [p_\phi, p_\lambda, p_h]^\top$ are the geographical coordinates (p_ϕ and p_λ are expressed in radian and p_h in meter), and ν is a white noise. The gravimeter observation equation (4.28) is expressed in mGal. The deterministic terms are:

- the nominal gravitational acceleration $g_{nom}(p_\lambda, p_h)$ calculated at the surface of the theoretical ellipsoid at latitude 45°. The following formula is used to calculate an approximate value of the acceleration of gravity at altitude p_h and at latitude p_λ.

$$g_{nom}(p_\lambda, p_h) \simeq 978031.846 \left(1 + 0.005278895 \sin^2(p_\lambda)\right) + 0.3086\ p_h \tag{4.29}$$

- the gravity anomalies map $map_{ga}(p_\phi, p_\lambda)$ that depends on the latitude and the longitude, is expressed in mGal.

- the vertical acceleration of the vehicle $\ddot{p}_h$ is derived from kinematic effects. Thereafter this term will be neglected as the AUV follows a uniform rectilinear trajectory.

- the Eötvös effect $Eötvös(V_E, p_\lambda)$ (also known as Coriolis force). It is an inertial force that acts on a moving vehicle within a reference frame that rotates (the Earth) with respect to an inertial frame. The centripetal acceleration that depends on the latitude is a part of the Eötvös effect but is ignored as it is negligible in comparison with the other terms of the observation equation (in particular with the gravity anomalies map output). The simplified Eötvös effect equation is given by:

$$Eötvös(V_E, p_\lambda) \simeq 3.7515\ V_E\ \cos(p_\lambda) \tag{4.30}$$

  where V_E is the velocity on the East axis expressed in meter per second.

Using the relations between the true position, the position given by the IMU, and the position errors (4.25), it comes:

$$Y^g = g_{nom}(\tilde{p}_\lambda+\delta p_\lambda, \tilde{p}_h+\delta p_h) + map_{ga}(\tilde{p}_\phi+\delta p_\phi, \tilde{p}_\lambda+\delta p_\lambda) + \ddot{\tilde{p}}_h + Eötvös(\tilde{V}_E+\delta V_E, \tilde{p}_\lambda+\delta p_\lambda) + \nu \tag{4.31}$$

where the East velocity $\tilde{V}_E$ is given by the IMU and the error δV_E is estimated by the filter. As for the position (4.25), the relations between the true velocity, the velocity given the IMU, and the position errors is: $\boldsymbol{V} = \tilde{\boldsymbol{V}} + \delta\boldsymbol{V}$.

Gravimeter Observation Equation for the Double Integrator Model

The observation equation for the double integrator model is:

$$Y^g = 7.66 \times 10^{-4}\ p_y - 0.3086\ p_z + 3.0542\ V_x + map_{ga}(p_x,\ p_y) + \nu \tag{4.32}$$

where $map_{ga}(p_x,\ p_y)$ is the simplified notation for $map_{ga}\left(\frac{p_x}{R\ cos(p_{\lambda 0})} + p_{\phi 0},\ \frac{p_y}{R} + p_{\lambda 0}\right)$, and ν is the measurements white noise of covariance $\boldsymbol{R}^g$. The measurement equation $\boldsymbol{Y}^g$ is expressed in mGal. The measurement equation (4.32) is derived from linearisations described in Appendix d.

The gravimeter-based navigation does not allow an accurate estimation of the position of the AUV because of the poor resolution of the gravity anomalies maps. However, it allows a better estimation of the velocity than the telemeter-based navigation as the velocity in the x-axis is directly observable in the observation equation (4.32).

4.2.3 *Sensor Fusion Architecture*

Bathymetric and gravimetric sensor fusion can be formulated in a centralised architecture, as illustrated in Figure 4.5. Centralized data fusion consists of designing an estimation filter able to gather information provided by the sensors and prior knowledge obtained from the dynamical and measurements models, and the IMU.

Performing data fusion from bathymetry and gravimetry yields several relevant advantages. Bathymetry observation equation depends on the vehicle position and on the seabed elevation profile. The multi-beam telemeter can thus provide information explicitly depending on position. However, velocity, which is not explicitly involved in the measurement model, can only be retrieved by the filter in an indirect way. The gravity observation equation depends on the geographical position and the x-axis velocity. The geographical position dependency brings additional information to bathymetric data. In addition, the explicit dependency of gravity observation equation on the x-axis velocity combined with the high accuracy of the atomic sensor brings a significant observability gain.

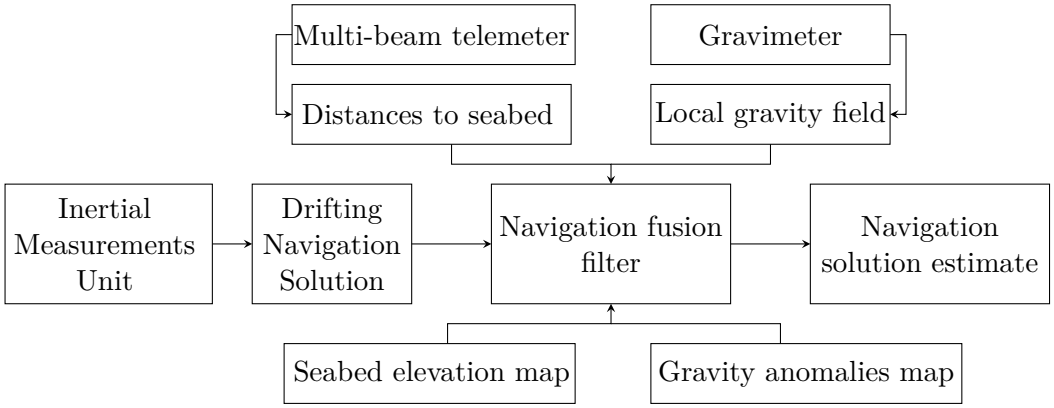

Figure 4.5: IMU hybridisation with a multi-beam telemeter and an atomic gravimeter in a centralised data fusion architecture.

Observability level can be quantified by the Fisher information matrix [52], which is the inverse of the Posterior Cramér-Rao Bound (PCRB) (see Appendix c). The deterministic recursive Tichavsky formulation of the Fisher information matrix, when the state model is linear and the measurement model is nonlinear, is:

$$\mathcal{I}(\boldsymbol{X}_{k+1}) = \mathbb{E}\left[\nabla h_{k+1}^{\top}(\boldsymbol{X}_{k+1})\, \boldsymbol{R}_{k+1}^{-1}\, \nabla h_{k+1}^{\top}(\boldsymbol{X}_{k+1})^{\top}\right] + \left(\boldsymbol{Q}_k + \boldsymbol{B}_k\, \mathcal{I}(\boldsymbol{X}_k)^{-1} \boldsymbol{B}_k^{\top}\right)^{-1} \tag{4.33}$$

where $\mathcal{I}(\boldsymbol{X}_k)$ is the prior Fisher information matrix and $\mathcal{I}(\boldsymbol{X}_{k+1})$ the posterior Fisher information matrix. See Appendix c for more details on the Fisher information matrix and the PCRB. The linear dynamical model is denoted $\boldsymbol{B}_k$ and $\boldsymbol{Q}_k$ is the process noise covariance. The two additive terms of the above equation (4.33) respectively quantify the measurement and the dynamics contributions to the Fisher information matrix. For both sensor configurations (bathymetry only, or gravimetry and bathymetry fusion) the dynamics contribution is the same and allows non explicitly observed variables (e.g. velocity) to be indirectly observed by integration correlation. Let us focus on the measurement contribution to the Fisher information matrix:

$$\mathcal{I}_m = \mathbb{E}\left[\nabla h_{k+1}^{\top}(\boldsymbol{X}_{k+1})\, \boldsymbol{R}_{k+1}^{-1}\, \nabla h_{k+1}^{\top}(\boldsymbol{X}_{k+1})^{\top}\right] \tag{4.34}$$

For the sake of simplicity, the computation of $\mathcal{I}_m$ is restricted to the double integrator formulations of the measurements equations. Also for simplicity, the bathymetry equation is approached as a single-beam telemeter pointing towards the local vertical direction:

$$Y^b = p_z - \mathrm{map}_{\mathrm{mb}}(p_x, p_y) + \nu \tag{4.35}$$

By applying (4.34), the bathymetry contribution to information can then be expressed as:

$$\mathcal{I}_m^b = \begin{bmatrix} \frac{(\nabla^x_{mb})^2}{\sigma^2_{mb}} & \frac{\nabla^x_{mb}\nabla^y_{mb}}{\sigma^2_{mb}} & -\frac{\nabla^x_{mb}}{\sigma^2_{mb}} & 0 & 0 & 0 \\ \frac{\nabla^x_{mb}\nabla^y_{mb}}{\sigma^2_{mb}} & \frac{(\nabla^y_{mb})^2}{\sigma^2_{mb}} & -\frac{\nabla^y_{mb}}{\sigma^2_{mb}} & 0 & 0 & 0 \\ -\frac{\nabla^x_{mb}}{\sigma^2_{mb}} & -\frac{\nabla^y_{mb}}{\sigma^2_{mb}} & \frac{1}{\sigma^2_{mb}} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \tag{4.36}$$

where $\nabla^x_{mb} = \frac{\partial \text{map}_{\text{mb}}(p_x,p_y)}{\partial p_x}$ and $\nabla^y_{mb} = \frac{\partial \text{map}_{\text{mb}}(p_x,p_y)}{\partial p_y}$ are respectively the East and North gradients of the seabed elevation map, and σ_{mb} is the standard deviation of the telemeter measurement error.

By adding the gravity field measurement, the measurements equation becomes:

$$\boldsymbol{Y}^f = [\boldsymbol{Y}^b, Y^g]^\top \tag{4.37}$$

The contribution of both sensors to information can be expressed as:

$$\mathcal{I}_m^f = \begin{bmatrix} \frac{(\nabla^x_{mb})^2}{\sigma^2_{mb}} + \frac{(\nabla^x_{ga})^2}{\sigma^2_{ga}} & \frac{\nabla^x_{mb}\nabla^y_{mb}}{\sigma^2_{mb}} + \frac{\nabla^x_{ga}(\nabla^y_{ga}+a)}{\sigma^2_{ga}} & -\frac{\nabla^x_{mb}}{\sigma^2_{mb}} + \frac{\nabla^x_{ga}b}{\sigma^2_{ga}} & \frac{\nabla^x_{ga}c}{\sigma^2_{ga}} & 0 & 0 \\ \frac{\nabla^x_{mb}\nabla^y_{mb}}{\sigma^2_{mb}} + \frac{\nabla^x_{ga}(\nabla^y_{ga}+a)}{\sigma^2_{ga}} & \frac{(\nabla^y_{mb})^2}{\sigma^2_{mb}} + \frac{(\nabla^y_{ga}+a)^2}{\sigma^2_{ga}} & -\frac{\nabla^y_{mb}}{\sigma^2_{mb}} + \frac{b(\nabla^y_{ga}+a)}{\sigma^2_{ga}} & \frac{c(\nabla^y_{ga}+a)}{\sigma^2_{ga}} & 0 & 0 \\ -\frac{\nabla^x_{mb}}{\sigma^2_{mb}} + \frac{\nabla^x_{ga}b}{\sigma^2_{ga}} & -\frac{\nabla^y_{mb}}{\sigma^2_{mb}} + \frac{b(\nabla^y_{ga}+a)}{\sigma^2_{ga}} & \frac{b^2}{\sigma^2_{ga}} + \frac{1}{\sigma^2_{mb}} & \frac{bc}{\sigma^2_{ga}} & 0 & 0 \\ \frac{\nabla^x_{ga}c}{\sigma^2_{ga}} & \frac{c(\nabla^y_{ga}+a)}{\sigma^2_{ga}} & \frac{bc}{\sigma^2_{ga}} & \frac{c^2}{\sigma^2_{ga}} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \tag{4.38}$$

where $\nabla^x_{ga} = \frac{\partial \text{map}_{\text{ga}}(p_x,p_y)}{\partial p_x}$ and $\nabla^y_{ga} = \frac{\partial \text{map}_{\text{ga}}(p_x,p_y)}{\partial p_y}$ are respectively the East and North gradients of the local gravimetry anomalies map, the constants a, b, and c are respectively equal to 7.66×10^{-4}, -0.3086, and 3.0542, and σ_{ga} is the standard deviation of the atomic gravimeter measurement error.

The matrix of the contribution of both sensors (4.38) shows, compared to equation (4.36), that gravimetry contributes to the observability of the following states:

- Position information is enhanced by complementary positive terms depending on the local gravity field gradient $(\nabla^x_{ga}, \nabla^y_{ga})^T$;
- x-axis velocity information becomes explicitly observable;
- Non diagonal terms bring cross information that will benefit to the whole state estimation.

All non-zero diagonal terms of the bathymetric Fisher information matrix are enhanced with positive (square) terms. These terms represent the theoretical impact of sensor fusion on expected estimation accuracy.

y-axis and z-axis velocities information remains the same, i.e., only depending on indirect estimation via information brought by the dynamics. A significant improvement is then expected on position estimation and x-axis velocity estimation. The estimation error reduction of these variable may also benefit to the others velocities estimation accuracy. Furthermore, atomic gravimeter offers a very low level noise σ_{ga} [10], which increases the impact of the gravimeter's contribution to information (proportional to $1/\sigma_{ga}^2$).

The complete Fisher information matrices for the multi-beam telemeter and both sensors are provided in Appendix e. To prove the contribution of the atomic gravimeter, the determinants of the complete Fisher information matrices are derived by taking the assumptions described in Appendix e. The determinant of the Fisher information matrix in the bathymetric setting (e.4) can be expressed as:

$$\det(\mathcal{I}^b) = \frac{\alpha_{mb}\,\Delta t^2 + \beta_{mb}}{q_{mb}} \tag{4.39}$$

where the constants α_{mb}, β_{mb}, and q_{mb} are given in Appendix e. The determinant of the Fisher information matrix when both sensors are activated (e.5) is:

$$\det(\mathcal{I}^f) = \frac{\alpha_{ga}\,\Delta t^2 + \beta_{ga}\Delta t + \gamma_{ga} + o(\Delta_t^2, {\nabla_{mb}^x}^2, {\nabla_{ga}^x}^2)}{q_{ga}} \tag{4.40}$$

where the constants α_{ga}, β_{ga}, γ_{ga} and q_{ga} are given in Appendix e. Any quadratic product of the gradient of the maps of order greater than or equal to 3 were removed of the determinant in equation (4.40). For several numerical values taken from our simulations settings (see Tables 5.4 or 6.1), we have:

$$\frac{\det(\mathcal{I}^f)}{\det(\mathcal{I}^b)} > 1 \tag{4.41}$$

As the determinant of the information matrix when the gravimeter is activated is greater than the determinant of the information matrix with only the multi-beam telemeter, the additional sensor brings more informations and thus leads to more accurate estimates.

4.3 TERRAIN AIDED NAVIGATION SCENARIO

Estimation algorithms are tested in a terrain scenario. The chosen area is the Californian coast (35°51′ N, 121°27′ W) shown in Figure 4.6.

The bathymetric and gravity anomalies maps are respectively plotted in Figures 4.7 and 4.8. The bathymetric map is available on the *Naval Postgraduate School* website [63]. The resolution is 200 m. The maximum depth is 4 km.

The gravity anomalies map is available on the *International Gravimetric Bureau* website [42]. The resolution is 2 km.

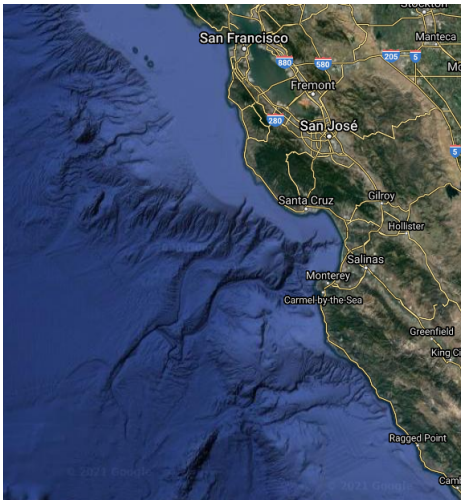

Figure 4.6: Californian coast (35°51′ N, 121°27′ W).

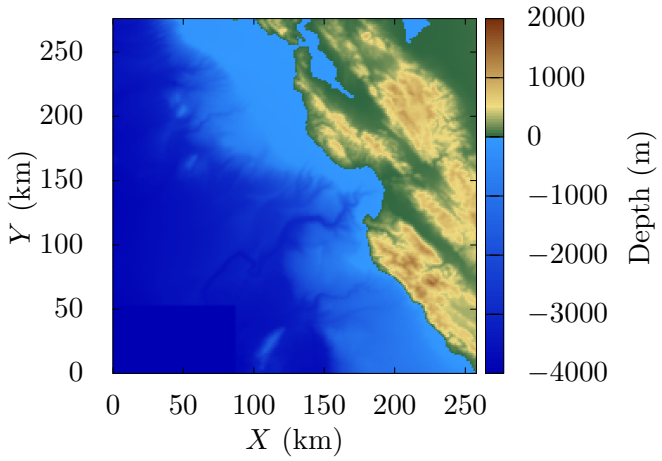

Figure 4.7: Bathymetric map of the Californian coast (35°51′ N, 121°27′ W) converted in Cartesian coordinates. The colorbar represents the depth levels (m).

The selected trajectory is shown in Figure 4.9. It is a rectilinear uniform trajectory of 10 minutes. The initial position of the trajectory is $[111000, 140000, -100]^\top$ m. The initial velocity of the AUV is $[5, 5, 0.05]^\top$ m s^{-1}.

The terrain profiles covered by the AUV in Figure 4.7 are shown in Figure 4.10. In the Figure 4.10a, the profile for the bathymetric map is created for a single-beam telemeter pointing towards the local vertical direction.

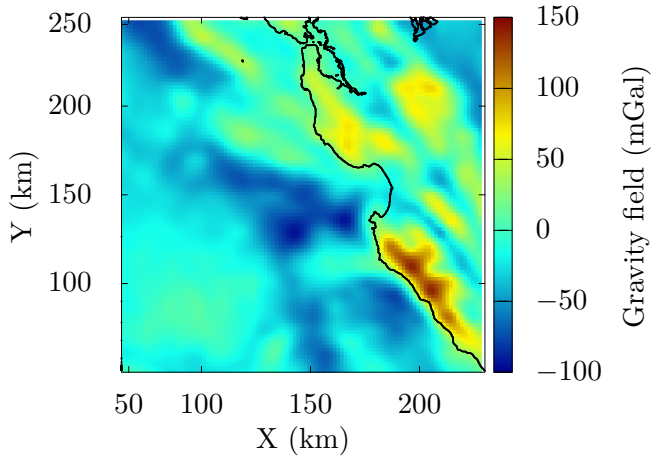

Figure 4.8: Gravimetric map of the Californian coast (35°51′ N, 121°27′ W) converted in Cartesian coordinates. The colorbar represents the gravity anomalies in (mGal).

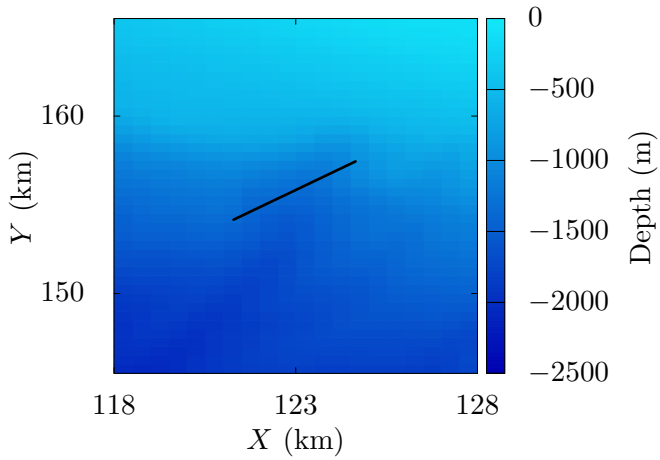

Figure 4.9: Bathymetric map of the area covered by the AUV trajectory.

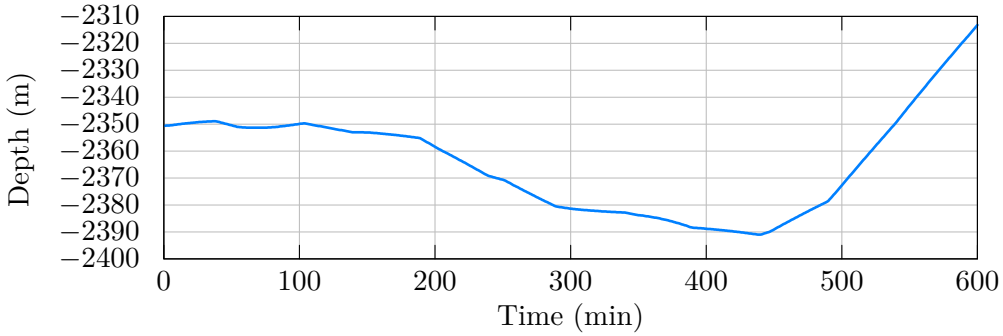

(a) Bathymetric profile.

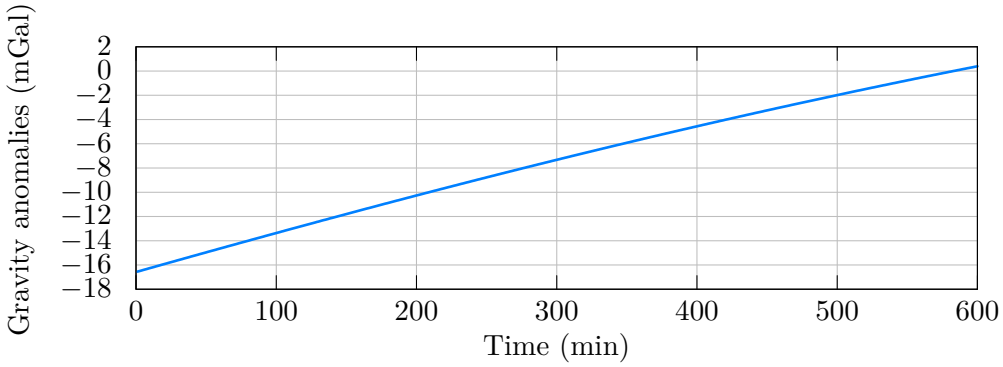

(b) Gravimetric profile.

Figure 4.10: Terrain profiles covered by the AUV.

4.4 EVALUATION CRITERIA FOR STATE ESTIMATION

In the following chapters, algorithms will be evaluated with several criteria defined in this section. Estimation algorithms outputs are a state estimate and a state confidence. In this manuscript, the state estimate is the sample mean $\hat{\boldsymbol{X}}_k$ and the state confidence is the covariance $\boldsymbol{P}_k$ of the approximate conditional density. Estimators are evaluated for $N_{MC} \in \mathbb{N}^*$ Monte Carlo simulations.

4.4.1 *Root Mean Square Error*

The accuracy of an algorithm is evaluated by the Root Mean Square Error (RMSE) defined by:

$$\mathrm{RMSE}_k^{\boldsymbol{\varkappa}} = \sqrt{\frac{1}{N_{MC}} \sum_{i=1}^{N_{MC}} \left\| \hat{\boldsymbol{\varkappa}}_k^i - \boldsymbol{\varkappa}_k^i \right\|^2} \tag{4.42}$$

where $\boldsymbol{\varkappa}_k^i$ and $\hat{\boldsymbol{\varkappa}}_k^i$ are respectively the actual state and the estimate for simulation i at time-step k. The state $\boldsymbol{\varkappa}_k^i$ can be the whole state $\boldsymbol{X}_k$ or only a sub-vector, e.g., the position or the velocity. From this criterion is often derived the horizontal RMSE for the position or the velocity, which is defined by:

$$\mathrm{RMSE}_k^{\boldsymbol{\varkappa}_h} = \sqrt{\left(\mathrm{RMSE}_k^{\boldsymbol{\varkappa}_x}\right)^2 + \left(\mathrm{RMSE}_k^{\boldsymbol{\varkappa}_y}\right)^2} \tag{4.43}$$

4.4.2 *Computation of the Posterior Cramér-Rao Bound for the Terrain Aided Navigation Application in the Double Integrator Model*

The recursive equation of the PCRB is given in the Appendix c:

$$PCRB_{k+1} = \left(\mathcal{I}(\boldsymbol{x}_{k+1})\right)^{-1} \tag{4.44}$$

where the Fisher information matrix $\mathcal{I}(\boldsymbol{x}_{k+1})$ is given by:

$$\mathcal{I}(\boldsymbol{x}_{k+1}) = \mathbb{E}\left[\nabla h_{k+1}^\top(\boldsymbol{X}_{k+1})\, \boldsymbol{R}_{k+1}^{-1}\, \nabla h_{k+1}^\top(\boldsymbol{X}_{k+1})^\top\right] + \left(\boldsymbol{Q}_k + \boldsymbol{B}_k\, \mathcal{I}(\boldsymbol{x}_k)^{-1} \boldsymbol{B}_k^\top\right)^{-1} \tag{4.45}$$

In the double integrator model, the gradients matrices of the observation equation $\nabla h_{k+1}(\boldsymbol{X}_{k+1})$ for the multi-beam telemeter, and for the sensor fusion scenario are provided in this section. Unlike the developments on the Fisher information matrix performed in Section 4.2.3, the bathymetric sensor considered here can have multiple beams.

The multi-beam telemeter observation equation can be redefined using the plane-line intersection formula:

$$r_i = \frac{-\boldsymbol{n}^\top \boldsymbol{p} + \boldsymbol{n}^\top \boldsymbol{p}^i}{\boldsymbol{n}^\top \boldsymbol{u}} \tag{4.46}$$

where $\boldsymbol{p}$ is the vehicle position, $\boldsymbol{p}^i$ the beam intersection point, $\boldsymbol{n}$ the terrain normal at the intersection point, and $\boldsymbol{u}$ the unit direction vector of the beam in the inertial frame (see Meduna [57]). The terrain normal is given by the following vector:

$$\boldsymbol{n} = \left[\nabla^x_{\text{mb}}, \nabla^y_{\text{mb}}, -1\right]^\top \tag{4.47}$$

where ∇^x_{mb} and ∇^y_{mb} are respectively the East and North gradients of the seabed elevation map. The gradient of the multi-beam telemeter observation equation (4.46) is:

$$\nabla h^b_k(\boldsymbol{X}_k) = \left[\frac{-\boldsymbol{n}^\top}{\boldsymbol{n}^\top \boldsymbol{u}}, \boldsymbol{0}_3\right] \tag{4.48}$$

The gravimeter observation equation for the double integrator model (4.32) is given by:

$$Y^g = 7.66 \times 10^{-4}\ p_y - 0.3086\ p_z + 3.0542\ V_x + map_{ga}(p_x,\ p_y) + \nu \tag{4.49}$$

The gradient of the gravimeter observation equation (4.32) is:

$$\nabla h^g_k(\boldsymbol{X}_k) = \left[\nabla^x_{\text{ga}}, a + \nabla^y_{\text{ga}}, -b, c, \boldsymbol{0}_2\right] \tag{4.50}$$

where the constants a, b, and c are respectively equal to 7.66×10^{-4}, -0.3086, and 3.0542.

The gradient of the observation equation $\nabla h_k(\boldsymbol{X}_k)$ is equal to $\nabla h^b_k(\boldsymbol{X}_k)$ in the bathymetric scenario. When the multi-beam telemeter and the gravimeter are activated, the gradient of the observation equation becomes:

$$\nabla h_k(\boldsymbol{X}_k) = \left[\nabla h^b_k(\boldsymbol{X}_k), \nabla h^g_k(\boldsymbol{X}_k)\right]^\top \tag{4.51}$$

The covariance of the observation process $\boldsymbol{R}_k \in \mathbb{R}^{(m+1)\times(m+1)}$ where m is the number of beams. All terms of $\boldsymbol{R}_k$ outside the diagonal are zero. The diagonal is equal to $[\sigma^2_{mb,1}, \ldots, \sigma^2_{mb,m}, \sigma^2_{ga}]$ where $\sigma_{mb,i}$ is the uncertainties associated with the i-th beam of the telemeter.

4.4.3 *Non-Convergence Rate*

The filter is qualified as non-convergent if its estimate $\hat{\boldsymbol{X}}_k$ leaves the confidence ellipsoid Γ_k, during the last five consecutive time-steps. In other terms, the filter is said to be non-convergent if:

$$\sum_{k=K-4}^{K} \mathbb{1}_{\Gamma_k}(\hat{\boldsymbol{X}}_k) = 0 \tag{4.52}$$

where K is the last time-step, and the ellipsoid Γ_k, given by the state confidence matrix $\boldsymbol{P}_k$, is defined by:

$$\Gamma_k(\boldsymbol{P}_k) = \left\{\boldsymbol{X}_k \in \mathbb{R}^d | (\boldsymbol{X}_k - \hat{\boldsymbol{X}}_k)^\top {\boldsymbol{P}_k}^{-1} (\boldsymbol{X}_k - \hat{\boldsymbol{X}}_k) \leq \alpha^2\right\} \tag{4.53}$$

The threshold α is given by $\mathbb{P}(\chi^2(d) \leq \alpha^2) = 0.99$ where d is the state dimension and χ^2 the Chi-squared distribution. In practice, the state confidence is either the covariance $\boldsymbol{P}_k$ given by the filter or the PCRB.

The non-convergence rate is the number of non-convergent Monte Carlo simulations out of the total number of simulation:

$$\frac{100}{N_{MC}} \sum_{i=1}^{N_{MC}} \mathbb{1}_{\left\{\sum_{k=K-4}^{K} \mathbb{1}_{\Gamma_k}(\hat{\boldsymbol{X}}_k^i)=0\right\}} \tag{4.54}$$

5

ADAPTIVE APPROXIMATE BAYESIAN COMPUTATION FILTERS

The consistency between the likelihood and the prior density is crucial to avoid degeneracy problem and filter divergence. Particle filters presented in Chapter 3 (e.g., the Regularised Particle Filter (RPF) and the Rao-Blackwellised Particle Filter (RBPF)) solve the filtering problem in the non-linear non-Gaussian case. However, standard Bayesian filtering algorithms do not guarantee the consistency between the likelihood and the prior density and assume an a priori knowledge on the measurement noise law. The Approximate Bayesian Computation (ABC) method was introduced (see Section 3.4) to cope with unknown likelihood. ABC consists in approximating the likelihood by a parametric kernel density function. The bandwidth parameter of this kernel is difficult to tune, which may yield filter instability.

This chapter introduces several contributions to likelihood-free particle filters:

- Section 5.1 introduces a new adaptive metric for choosing the bandwidth parameter of the kernel density function, which is based on the effective sample size criterion and is called the Adaptive Approximate Bayesian Computation (A2BC) method.
- Section 5.2 studies the impact of the A2BC on two error quantifiers: the asymptotic variance of the unnormalised weights and the local Monte Carlo error.

A computational load study is provided in Section 5.2.3, in terms of theoretical floating-point operations per time-step (see Appendix f). Section 5.3 compares the performances in terms of RMSE of the A2BC-RPF and the A2BC-RBPF with their conventional versions (i.e., RPF and RBPF) on a TAN example.

5.1 PRINCIPLE

Consider the state-space model given in Chapter 2 by equations (2.17) and (2.2):

$$\begin{cases} \boldsymbol{X}_k = b_k(\boldsymbol{X}_{k-1}) + \boldsymbol{W}_k \\ \boldsymbol{Y}_k = h_k(\boldsymbol{X}_k) + \boldsymbol{V}_k \end{cases} \tag{5.1}$$

where $b_k : \mathbb{R}^d \to \mathbb{R}^d$ and $h_k : \mathbb{R}^d \to \mathbb{R}^{d_y}$ are respectively the dynamical and observation functions, $\boldsymbol{W}_k$ and $\boldsymbol{V}_k$ are the process and observation white noises. The initial state $\boldsymbol{X}_0$ of

known density $p(\boldsymbol{X}_0)$ is independent of both noises. The observations $\{\boldsymbol{Y}_k\}_{k\geq 0}$ are mutually independent conditionally to the state $\{\boldsymbol{X}_k\}_{k\geq 0}$. The noises $\boldsymbol{W}_k$ and $\boldsymbol{V}_k$ are assumed mutually independent.

The law of the observation noise $\boldsymbol{V}_k$ is assumed to be unknown. This is often the case in TAN applications (see Section 4.2) that use embedded numerical terrain maps, which can lead to uncontrolled observation errors. Maps are not exempt from uncertainty, they may not be fully representative of the field that a local sensor would measure. Furthermore, the multi-beam telemeter observation model cannot be put into an explicit form that would link the state to the observations. This is true for both the inertial errors model (4.26) and the double integrator model (4.27). The multi-beam telemeter observation model is then numerically approximated which can lead to uncontrolled errors. In this context, neither the observation noise (involving both the sensor noise and the map modeling error), nor its probabilistic law (the likelihood $p(\boldsymbol{Y}_k|\boldsymbol{X}_k)$) can be analytically determined. Only the accuracy of involved sensors are known in the form of a covariance $\boldsymbol{R}_k \in \mathbb{R}^{d_y \times d_y}$.

In order to tackle the issue of unknown likelihood, we propose to replace it by a probability density kernel with an adaptive bandwidth parameter. Similarly to ABC method (see Section 3.4), the approximate likelihood $p(\boldsymbol{Y}_k|\boldsymbol{X}_k)$ is given by:

$$p^{\varepsilon_k}(\boldsymbol{Y}_k|\boldsymbol{X}_k) = \mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - \boldsymbol{U}_k) \tag{5.2}$$

where the kernel density $\mathcal{K}_{\varepsilon_k} : \mathbb{R}^{d_y} \to \mathbb{R}$ is set by its bandwidth parameter $\varepsilon_k \geq 1$. The pseudo-observations $\boldsymbol{U}_k$ are derived from the noise-free observation model:

$$\boldsymbol{U}_k = h_k(\boldsymbol{X}_k) \tag{5.3}$$

The kernel density $\mathcal{K}_{\varepsilon_k}$ takes into account the known measurement noise covariance $\boldsymbol{R}_k$. For example, the likelihood can be approximated by a Gaussian kernel:

$$\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - \boldsymbol{U}_k) \propto \exp\left(\frac{-1}{2}(\boldsymbol{Y}_k - \boldsymbol{U}_k)^\top (\varepsilon_k^2 \boldsymbol{R}_k)^{-1}(\boldsymbol{Y}_k - \boldsymbol{U}_k)\right) \tag{5.4}$$

or by a Cauchy kernel:

$$\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - \boldsymbol{U}_k) \propto \left(1 + (\boldsymbol{Y}_k - \boldsymbol{U}_k)^\top (\varepsilon_k^2 \boldsymbol{R}_k)^{-1}(\boldsymbol{Y}_k - \boldsymbol{U}_k)\right)^{-\frac{1+d_y}{2}} \tag{5.5}$$

Using the kernel density as likelihood, the optimal filter (see Section 2.3) becomes:

- The prediction step determines the prior density $p(\boldsymbol{X}_k|\mathcal{Y}_k)$ via the Chapman-Kolmogorov equation (2.20):

$$p(\boldsymbol{X}_k|\mathcal{Y}_{k-1}) = \int p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1})\; p(\boldsymbol{X}_{k-1}|\mathcal{Y}_{k-1})\; d\boldsymbol{X}_{k-1} \tag{5.6}$$

- From the Bayes' law, the correction step is derived:

$$p(\boldsymbol{X}_k|\mathcal{Y}_k) = \frac{\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - \boldsymbol{U}_k)\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{\int \mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - \boldsymbol{U}_k)\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\; d\boldsymbol{X}_k} \tag{5.7}$$

5.1.1 *Choice of the Bandwidth Parameter*

The bandwidth parameter ε_k controls the support of the likelihood and has thus a strong impact on the correction step of particle filters (see Section 3.3.1). The performance of Monte Carlo methods depends on the consistency between the likelihood and the prior density (see Section 3.5). An unsuitable choice of likelihood can lead to an inefficient correction step (see Section 3.3.1, Figure 3.2).

The choice of the bandwidth parameter has a strong impact on the asymptotic variance of the unnormalised weights (3.56) as the likelihood is replaced by the kernel density $\mathcal{K}_{\varepsilon_k}$. When the proposal density is equal to the prior density, the approximate asymptotic variance of the unnormalised weights is given by:

$$\mathbb{V}_{\varepsilon_k} = \frac{\int \mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))^2 \, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1}) \, d\boldsymbol{X}_k}{\left(\int \mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k)) \, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1}) \, d\boldsymbol{X}_k\right)^2} - 1 \tag{5.8}$$

To illustrate the influence of the bandwidth parameter on the approximate asymptotic variance of the unnormalised weights, only in this paragraph, the following assumptions are taken:

- The dimensions of the state and observations are equal to $d = d_y = 1$;
- The prior density is a Gaussian:

$$p(X_k|\mathcal{Y}_{k-1}) = \frac{1}{\sqrt{2\pi}\,\sigma_{\boldsymbol{P}}} \exp\left(\frac{-(X_k - \tilde{X}_k)^2}{2\,\sigma_{\boldsymbol{P}}^2}\right) \tag{5.9}$$

  where $\tilde{X}_k$ and $\sigma_{\boldsymbol{P}}$ are respectively the predicted state and standard deviation;

- The kernel density is Gaussian:

$$\mathcal{K}_{\varepsilon_k}(Y_k - h_k(X_k)) = \frac{1}{\sqrt{2\pi}\,\varepsilon_k\,\sigma_{\boldsymbol{R}}} \exp\left(\frac{-(Y_k - h_k(X_k))^2}{2(\varepsilon_k\,\sigma_{\boldsymbol{R}})^2}\right) \tag{5.10}$$

  where $\sigma_{\boldsymbol{R}}$ is the standard deviation of the measurement noise;

- The pseudo-measurements are linearised around the estimate $\tilde{X}_k$: $h_k(X_k) \approx h_k(\tilde{X}_k) + (X_k - \tilde{X}_k)\nabla h_k(\tilde{X}_k)$.

The following developments are similar in the multidimensional case. Using the Gaussian assumptions (5.9)-(5.10), it comes:

$$\int \mathcal{K}_{\varepsilon_k}(Y_k - h_k(X_k))^2 \, p(X_k|\mathcal{Y}_{k-1}) \, dX_k = \frac{1}{2\pi\varepsilon_k((\varepsilon_k\,\sigma_{\boldsymbol{R}})^2 + 2\nabla h_k(\tilde{X}_k)^2\sigma_{\boldsymbol{P}}^2)^{\frac{1}{2}}} \exp\left(-\frac{(Y_k - h_k(\tilde{X}_k))^2}{(\varepsilon_k\,\sigma_{\boldsymbol{R}})^2 + 2\nabla h_k(\tilde{X}_k)^2\sigma_{\boldsymbol{P}}^2}\right) \tag{5.11}$$

$$\int \mathcal{K}_{\varepsilon_k}(Y_k - h_k(X_k))\, p(X_k|\mathcal{Y}_{k-1})\, dX_k = \frac{1}{\sqrt{2\pi}((\varepsilon_k\, \sigma_{\boldsymbol{R}})^2 + \nabla h_k(\tilde{X}_k)^2 \sigma_{\boldsymbol{P}}^2)^{\frac{1}{2}}} \exp\left(-\frac{(Y_k - h_k(\tilde{X}_k))^2}{2((\varepsilon_k\, \sigma_{\boldsymbol{R}})^2 + \nabla h_k(\tilde{X}_k)^2 \sigma_{\boldsymbol{P}}^2)}\right) \tag{5.12}$$

Using equations (5.11) and (5.12), the approximate asymptotic variance of the unnormalised weights is:

$$\mathbb{V}_{\varepsilon_k} = \frac{(\varepsilon_k\, \sigma_{\boldsymbol{R}})^2 + \nabla h_k(\tilde{X}_k)^2 \sigma_{\boldsymbol{P}}^2}{\varepsilon_k((\varepsilon_k\, \sigma_{\boldsymbol{R}})^2 + 2\nabla h_k(\tilde{X}_k)^2 \sigma_{\boldsymbol{P}}^2)^{\frac{1}{2}}} \exp\left(\boldsymbol{C}(Y_k - h_k(\tilde{X}_k))^2\right) - 1 \tag{5.13}$$

where

$$\boldsymbol{C} = \frac{\nabla h_k(\tilde{X}_k)^2 \sigma_{\boldsymbol{P}}^2}{((\varepsilon_k\, \sigma_{\boldsymbol{R}})^2 + 2\nabla h_k(\tilde{X}_k)^2 \sigma_{\boldsymbol{P}}^2)((\varepsilon_k\, \sigma_{\boldsymbol{R}})^2 + \nabla h_k(\tilde{X}_k)^2 \sigma_{\boldsymbol{P}}^2)} \tag{5.14}$$

From equation (5.13), the following points are concluded:

- The approximate asymptotic variance tends toward infinity when the bandwidth parameter ε_k tends to 0.
- When the parameter ε_k tends towards infinity, the approximate asymptotic variance tends towards 0.

These results illustrate the importance of the bandwidth parameter on the approximate asymptotic variance of the unnormalised weights.

As demonstrate in Appendix b, the effective sample size criterion (3.18) can be approximated by an expression that depends on the asymptotic variance of the unnormalised weights:

$$\text{ESS} = \frac{N}{1 + \mathbb{V}\text{ar}_q(\tilde{w})} \approx \frac{N}{1 + \mathbb{V}} \tag{5.15}$$

As a result, when the asymptotic variance decreases, the effective sample size criterion increases. Thus, the consistency between the likelihood and the prior density, and the effective sample size increase simultaneously.

The effective sample size represents the number of effective particles, i.e., particles which have a significant weight and contribute to the estimation of the conditional density. The idea behind the proposed method is to choose ε_k such that the number of effective particles is greater than a given resampling threshold $N_{\text{th}} = \theta_{\text{th}}\, N$ where $\theta_{\text{th}} \in [0, 1]$ and N is the total number of particles. From the previous paragraph, we know that finding ε_k such that the effective sample size increases is equivalent to increasing the consistency between the likelihood and the prior density.

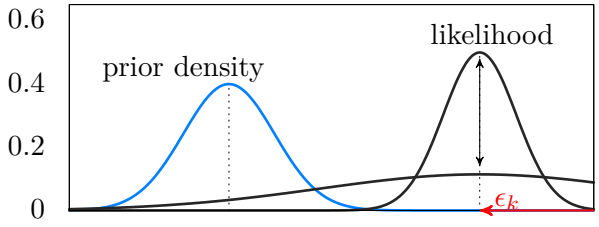

Figure 5.1: Consistency between the prior density and a flattened non-informative likelihood ($\varepsilon_k \to \infty$) resulting in inefficient correction of the particles.

This choice of bandwidth parameter avoids the resampling. However, choosing ε_k such that the number of effective particles is greater than N_{th} can lead to an inefficient correction of the particles, as illustrated in Figure 5.1. When the particles are significantly misplaced in the state space, the support of the prior density does not overlap with the support of the likelihood. This is the case in Figure 5.1 between the prior density and the sharp likelihood, which corresponds to the case where $\varepsilon_k = 1$. In this situation, the support of the likelihood is increased (i.e., $1 < \varepsilon_k \to \infty$) until the consistency between the likelihood and the prior density guarantees N_{th} effective particles. By increasing the support of the likelihood, the information brought by observations is lost. Indeed, the observations no longer discriminate the particles significantly misplaced by attributing to them a weight close to zero, and a majority of particles is blindly considered effective. The correction step is thus ineffective and the filter may diverge. In oder to access the resampling when necessary, the bandwidth parameter ε_k is restricted to a domain.

The objective of the proposed method, called A2BC, is then to optimally increase the support of the likelihood to reach N_{th} effective particles without losing the information brought by sensor observations. Formally, we resolve:

$$\varepsilon_k^\star = \underset{\varepsilon_k \in \mathcal{D}}{\operatorname{argmin}} \left(\text{ESS}(\varepsilon_k) - N_{\text{th}}\right)^2 \tag{5.16}$$

where $\mathcal{D} \subset \mathbb{R}^+$ is a given domain. The notation are slightly modified to reflect the dependency with the bandwidth parameter ε_k:

$$\text{ESS}(\varepsilon_k) = \frac{N}{1 + \mathbb{V}\text{ar}_q(\tilde{w}_{\varepsilon_k})} \tag{5.17a}$$

$$\tilde{w}_{\varepsilon_k} = \frac{\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \tag{5.17b}$$

The optimisation problem (5.16) is restricted to the domain $\mathcal{D}$ to access the resampling step whenever the particles are significantly misplaced. As the occurrence of the resampling step is reduced, the A2BC method avoids the errors associated to it. Indeed, resampling can cause the premature loss of one or more modes of the conditional density. When the lost mode corresponds to the true position, the filter automatically diverges. By delaying the

resampling, the convergence to a single mode of the likelihood is delayed, which makes the filter more robust to measurements ambiguities.

The bandwidth parameter ε_k belongs to the domain $\mathcal{D} = [\varepsilon_k^-, \varepsilon_k^+]$. The coefficient must be at least equal to unity (i.e., $\mathcal{D} = [1, \varepsilon_k^+]$) as the observation errors are at least equal to the sensor errors. The upper bound of the bandwidth ε_k^+ is chosen according to the application and must not allow the likelihood to be flattened.

The choice of the parameter ε_k (5.16) is summarised by the three following situations:

- $\text{ESS}(1) > N_{\text{th}}$: If the number of effective particles is greater than the resampling threshold by taking $\varepsilon_k = 1$, then the chosen bandwidth parameter will be $\varepsilon_k^\star = 1$. In this situation, increasing the support of the likelihood is not necessary and the information provided by observations is preserved.
- $\text{ESS}(1) < N_{\text{th}} < \text{ESS}(\varepsilon_k)$: In this case, the bandwidth ε_k^* is chosen such that $\text{ESS}(\varepsilon_k^\star) - N_{\text{th}} = 0$.
- $\text{ESS}(\varepsilon_k^+) < N_{\text{th}}$: If the number of effective particles calculated with $\varepsilon_k = \varepsilon_k^+$ is below the threshold N_{th}, then the chosen bandwidth parameter (that satisfies the optimisation problem (5.16)) is $\varepsilon_k^\star = \varepsilon_k^+$ and the resampling step is triggered.

The proposed choice of bandwidth parameter (5.16) has not a direct impact on the first and third cases. It is particularly suitable when the number of effective particles is close to the threshold N_{th} chosen by the user. In these critical moments, resampling may cause the filter to diverge by not selecting the correct mode of the conditional density. Limiting the resampling makes it possible to wait for the measurements to remove the ambiguities. The proposed bandwidth parameter improves the filter's robustness to measurements uncertainties. This is crucial for underwater navigation as, in the TAN method, the underwater maps are ambiguous (flat areas or areas with similar characteristics). The objective of the proposed method is to prevent the filter from diverging by optimising the impact of the weights correction.

As the effective sample size cannot be evaluated exactly, it is replace in practice by its estimation (see Section 3.3.1 equation (3.19)):

$$\widehat{\text{ESS}}(\varepsilon_k) = \frac{1}{\sum_{i=1}^{N} \left(w_{\varepsilon_k}^i\right)^2} \tag{5.18}$$

The optimisation problem becomes:

$$\varepsilon_k^\star = \underset{\varepsilon_k \in \mathcal{D}}{\text{argmin}} \left(\widehat{\text{ESS}}(\varepsilon_k) - N_{\text{th}}\right)^2 \tag{5.19}$$

5.1.2 *Adaptive Approximate Bayesian Computation method within Particle Filters*

The A2BC method described in Section 5.1 can be integrated within any particle filter. Only the correction step is modified. A2BC particle filters also stem from the Sequential Importance

Algorithm 5.1 A2BC-SIR-PF

Initialisation: The initial particle set is drawn as $\{\boldsymbol{X}_0^i\}_{i\in[1,N]}$ using the initial density $p(\boldsymbol{X}_0)$ and the initial weights set $\{w_0^i\}_{i\in[1,N]}$ is taken equal to $\frac{1}{N}$.
for each time-step k **do**
  Prediction: Draw the particles from a suitable proposal density $\boldsymbol{X}_k^i \sim q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}^i, \boldsymbol{Y}_k)$.
  Simulate the pseudo-measurements $\boldsymbol{U}_k^i$ by plugging the particles $\boldsymbol{X}_k^i$ in the observation model h_k: $\boldsymbol{U}_k^i = h_k(\boldsymbol{X}_k^i)$.
  Correction: Resolve the optimisation problem (5.16) to determine the bandwidth parameter $\epsilon_k^\star$:

$$\epsilon_k^\star = \underset{\epsilon_k\in\mathcal{D}}{\operatorname{argmin}} \left(\widehat{\mathrm{ESS}}(\epsilon_k) - N_{\mathrm{th}}\right)^2 \tag{5.20}$$

  Update the weights using the A2BC likelihood:

$$w_k^i \propto w_{k-1}^i \frac{\mathcal{K}_{\epsilon_k^\star}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k^i))\; p(\boldsymbol{X}_k^i|\boldsymbol{X}_{k-1}^i)}{q(\boldsymbol{X}_k^i|\boldsymbol{X}_{k-1}^i, \boldsymbol{Y}_k)} \tag{5.21}$$

  Compute the state estimate $\hat{\boldsymbol{X}}_k$ (3.16) and its covariance $\boldsymbol{P}_k$ (3.17).
  if $\widehat{ESS} < N_{\mathrm{th}}$ **then**
    Draw a new set of particles $\{\boldsymbol{X}_k^i\}_{i\in[1,N]}$ using a resampling method, e.g., the multinomial resampling (Algorithm 3.2).
    Reset the weights to $\frac{1}{N}$.
  end if
end for
Return the state estimate $\hat{\boldsymbol{X}}_k$ and its covariance $\boldsymbol{P}_k$, $\forall k$.

Resampling Particle Filter (SIR-PF) algorithm presented in Section 3.3. The A2BC-SIR-PF is summarised in Algorithm 5.1.

Two A2BC filters are derived from the conventional RPF and RBPF (see Section 3.3): the A2BC-RPF and the A2BC-RBPF respectively. Both filters are summarised in Appendix g. The particle filter used in the RBPF is the RPF.

5.2 THEORETICAL RESULTS

The Monte Carlo filters performance criteria presented in Section 3.5 are driven by the consistency between the likelihood and the prior density. The consistency is ensured by the A2BC method, which therefore has a direct impact on the performance of the Monte Carlo filters.

Section 5.2.1 studies the impact of the A2BC method on the approximate asymptotic variance of the unnormalised weights. Section 5.2.2 presents its impact on the local Monte Carlo error. The algorithmic complexity of the A2BC method is discussed in Section 5.2.3 to evaluate the computation cost of such an approach in an industrial embedded context.

5.2.1 *Upper Bound of the Variance of the Unnormalized Weights*

Proposition 2 states that the approximate asymptotic variance of the unnormalised weights is bounded if the discrepancy remains below a given threshold. The discrepancy (3.63) γ_{ε_k} is computed using the weights $\tilde{w}_{\varepsilon_k}$ (5.17b). A small variance of the unnormalised weights guarantees an accurate Monte Carlo estimate (see Section 3.1).

Proposition 2. *By choosing ε_k such that $\gamma_{\varepsilon_k} \leq \gamma_{th}$, where γ_{th} is a threshold, it is guaranteed that the asymptotic variance of the unnormalised weights $\mathbb{V}_{\varepsilon_k}$ is bounded:*

$$\mathbb{V}_{\varepsilon_k} < +\infty \tag{5.22}$$

Proof. By taking the supremum, it comes:

$$\int \frac{\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))^2\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})^2}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)}\, d\boldsymbol{X}_k \leq \sup_{\boldsymbol{X}_k \in \mathbb{R}^d} \left(\frac{\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \right) \int \mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k \tag{5.23}$$

It follows that:

$$\mathbb{V}_{\varepsilon_k} = \frac{\int \frac{\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))^2\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})^2}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)}\, d\boldsymbol{X}_k}{\left(\int \mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k\right)^2} - 1 \tag{5.24a}$$

$$\leq \frac{\sup_{\boldsymbol{X} \in \mathbb{R}^d} \left(\frac{\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \right)}{\int \mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k} - 1 = \gamma_{\varepsilon_k} - 1 \tag{5.24b}$$

By choosing ε_k such that $\gamma_{\varepsilon_k} \leq \gamma_{\text{th}}$, where γ_{th} is a threshold, it comes:

$$0 \leq \mathbb{V}_{\varepsilon_k} \leq \gamma_{\varepsilon_k} - 1 < +\infty \tag{5.25}$$

The variance admits an upper bound that depends on the discrepancy. The variance is small when the discrepancy γ_{ε_k} tends towards 1. □

From (3.18) and (5.25), and as $1 \leq \text{ESS}(\varepsilon_k) \leq N$, it comes:

$$\frac{N}{\gamma_{\varepsilon_k}} \leq \text{ESS}(\varepsilon_k) \leq N \tag{5.26}$$

When γ_{ε_k} tends towards 1, from the squeeze theorem, $\text{ESS}(\varepsilon_k)$ tends towards N. Therefore, when the discrepancy is small, a significant number of particles is efficient.

By taking N_{th} equal to $\frac{N}{\gamma_{\text{th}}}$, it follows:

$$N_{\text{th}} = \frac{N}{\gamma_{\text{th}}} \leq \frac{N}{\gamma_{\varepsilon_k}} \leq \text{ESS}(\varepsilon_k) \tag{5.27}$$

5.2.2 *Monte Carlo Local Error Analysis*

Proposition 3 shows that the control of the discrepancy γ_{ε_k} not only bounds the asymptotic variance of the unnormalised weights, but also reduces the local Monte Carlo error. The local Monte Carlo error quantifies the matching between the filter's density and the actual state density. It is defined by:

$$E_k \coloneqq \mathbb{E}[||\hat{\pi}_k - \pi_k||_1 \mid \mathcal{S}_{k-1}] \tag{5.28}$$

where π_k is the conditional distribution, $\hat{\pi}_k$ its approximation, $||.||_1$ the total variation norm, and $\mathcal{S}_{k-1}$ the σ-field generated by the measurements and the random variables simulated in the transition from $\hat{\pi}_0$ to $\hat{\pi}_k$.

Thereafter, $\pi_{k|k-1}$ denotes the predictive distribution and $\hat{\pi}_{k|k-1}$ its approximation.

Proposition 3. *If $\pi_{k|k-1}$ is absolutely continuous with respect to the Lebesgue measure and if the density $\frac{d\pi_{k|k-1}}{dx} \in W^{2,1}$, then the local Monte Carlo error verifies:*

$$E_k \leq \frac{\tilde{\alpha}}{N^{\frac{2}{d+4}}} + \gamma_{\varepsilon_k}\left[\frac{\tilde{\beta}}{N^{\frac{2}{d+4}}} + ||\hat{\pi}_{k|k-1} - \pi_{k|k-1}||_1\right] \tag{5.29}$$

The two constants $\tilde{\alpha} = \alpha(\mathcal{K}_h, \mathcal{K}_{\varepsilon_k}.\pi_{k|k-1})$ and $\tilde{\beta} = \beta(\mathcal{K}_h, \hat{\pi}_{k|k-1})$ depend on the parameters of the regularisation kernel $\mathcal{K}_h$, the A2BC kernel $\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))$, the predictive distribution $\pi_{k|k-1}$, and it approximation $\hat{\pi}_{k|k-1}$.

The following notations are used in the proof of Proposition 3:

- $W^{2,1}$ is the Sobolev space of measurable functions defined on $\mathbb{R}^d$ for which the derivatives up to order 2 (in the sense of the distributions) are in $L^1 \in \mathbb{R}$,
- $|.|_{2,1}$ is the usual norm on $W_{2,1}$,
- Let ν be a signed measure on E of density f with respect to Lebesgue's measurement, we define:
  - $a_2(\nu) = a_2(f) = 2!\ \max_{|\alpha|=2} \int_E |\boldsymbol{X}^\alpha||f(x)|dx$,
  - $A'_d = 2^{\frac{d}{2}} \left(\frac{(2\pi)^{\frac{d+1}{2}}}{\Gamma\left(\frac{d+1}{2}\right)} \right)^{\frac{1}{2}}$
  - $I(\nu) = I(f) = \int_E |\boldsymbol{X}|^{d+1} |f(\boldsymbol{X})| d\boldsymbol{X}$

Proof. The problem is stated with the following assumptions:

- The regularisation kernel $\mathcal{K}_h$ is radially symmetric and order 2 positive, which is the case for the Epanechnikov kernel;
- The likelihood is positive and bounded on $\mathbb{R}^d$ such that $\int \pi_k(d\boldsymbol{X}_k)\, p(\boldsymbol{Y}_k|\boldsymbol{X}_k) > 0$;

Under these assumptions, from the work of Oudjane [68] (Prop. 5.19 p.168) in the context of progressive correction, the local Monte Carlo error is bounded by:

$$E_k \leq \hat{\alpha} h^2 + \hat{\gamma}_k \left[\frac{\hat{\beta}}{\sqrt{N h^d}} + ||\hat{\pi}_{k|k-1} - \pi_{k|k-1}||_1 \right] \tag{5.30}$$

where $\hat{\alpha} = a_2(\mathcal{K}_h)\ |p(\boldsymbol{Y}_k|\boldsymbol{X}_k).\pi_{k|k-1}|_{2,1}$ and $\hat{\beta} = 2\ A'_d\ \left[1 + I(\hat{\pi}_{k|k-1})\right]\ ||\mathcal{K}_h||_2$ are two constants.

In the context of A2BC, the likelihood $p(\boldsymbol{Y}_k|\boldsymbol{X}_k)$ is replaced by $\mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))$, which depends on its bandwidth parameter ε_k. Let $\mathcal{K}_{\varepsilon_k}$ be positive and bounded on $\mathbb{R}^d$; assuming that $\varepsilon_k < \infty$, it follows:

$$E_k \leq \hat{\alpha} h^2 + \gamma_{\varepsilon_k} \left[\frac{\hat{\beta}}{\sqrt{N h^d}} + ||\hat{\pi}_{k|k-1} - \pi_{k|k-1}||_1 \right] \tag{5.31}$$

where $\hat{\alpha} = a_2(\mathcal{K}_h)|\mathcal{K}_{\varepsilon_k}.\pi_{k|k-1}|_{2,1}$ and $\hat{\beta} = 2A'_d \left[1 + I(\hat{\pi}_{k|k-1})\right] ||\mathcal{K}_h||_2$ are two constants.

As the regularisation kernel is the Epanechnikov kernel, it bandwidth parameter is given by equation (3.25): $h = \mu\ A(\mathcal{K}_h)\ N^{-\frac{1}{d+4}}$. By replacing the regularisation bandwidth parameter by its expression, it comes:

$$E_k \leq \frac{\tilde{\alpha}}{N^{\frac{2}{d+4}}} + \gamma_{\varepsilon_k} \left(\frac{\tilde{\beta}}{N^{\frac{2}{d+4}}} + ||\hat{\pi}_{k|k-1} - \pi_{k|k-1}||_1 \right) \tag{5.32}$$

where

$$\begin{cases} \tilde{\alpha} & = \mu^2 A(\mathcal{K}_h)^2 \hat{\alpha} \\ \tilde{\beta} & = \frac{\hat{\beta}}{\mu^{d/2} A(\mathcal{K}_h)^{d/2}} \end{cases} = \begin{cases} & \mu^2 A(\mathcal{K}_h)^2 a_2(\mathcal{K}_h) |\mathcal{K}_{\varepsilon_k} . \pi_{k|k-1}|_{2,1} \\ & \frac{2A'_d [1+I(\hat{\pi}_{k|k-1})] ||\mathcal{K}_h||_2}{\mu^{d/2} A(\mathcal{K}_h)^{d/2}} \end{cases} \tag{5.33}$$

□

The local Monte Carlo error E_k is the error after the correction step. E_k depends on the number of particles N, the prediction error $||\hat{\pi}_{k|k-1} - \pi_{k|k-1}||_1$, the discrepancy γ_{ε_k}, and the regularisation bandwidth parameter h. If the number of particles N is large, the two constant terms $\tilde{\alpha}$ and $\tilde{\beta}$ have little influence on the local Monte Carlo error as both terms $\frac{\tilde{\alpha}}{N^{\frac{2}{d+4}}}$ and $\frac{\tilde{\beta}}{N^{\frac{2}{d+4}}}$ tend towards zero. In this case, the local Monte Carlo error E_k depends on the coefficient γ_{ε_k} in front of the prediction error. The discrepancy γ_{ε_k} increases when the consistency between the prior density and the likelihood decreases. The A2BC method adjusts the likelihood by increasing slightly its support to ensure a better consistency with the prior density. As a result, the supremum of the likelihood decreases and the consistency with the prior density increases which guarantees that $\gamma_{\varepsilon_k} \leq \hat{\gamma}_k$.

To summarise, in Proposition 3, we proved that the local Monte Carlo error E_k admits an upper bound that depends on the discrepancy γ_{ε_k}. The use of the A2BC method in the correction step aims to limit the discrepancy by optimising the scale parameter of the likelihood. Therefore, an appropriate choice of the bandwidth parameter of the A2BC kernel ε_k makes it possible to optimise the upper bound of the local Monte Carlo error E_k.

5.2.3 *Complexity Analysis of the Choice of the Bandwidth Parameter*

Theoretical computational loads are given in terms of total number of floating-point operations (flops, see Appendix f). In what follows, the computation for the A2BC method consists of the choice of the bandwidth parameter of the A2BC kernel. The total load of the prediction, correction, resampling and regularisation steps are provided in Appendix f for the RPF and the RBPF filters and remain unchanged in the context of A2BC.

The computational load of the bandwidth parameter ε_k (5.16), for a fix k, is:

$$c_\varepsilon = c_{\mathcal{D}} \left(N(c_{\mathcal{K}} + c_p + c_q + 9) - 1 \right) \tag{5.34}$$

where $c_{\mathcal{K}}$ is the cost of the A2BC likelihood $p^{\varepsilon_k}(\boldsymbol{Y}_k|\boldsymbol{X}_k) = \mathcal{K}_{\varepsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k))$, c_p the cost of the prior density $p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})$ and c_q the cost of the proposal density $q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)$. The load of the determination of the bandwidth c_ε depends on the choice of the domain $\mathcal{D}$ and therefore on the number of calls of the objective function denoted $c_{\mathcal{D}}$.

To explain the computational load of the bandwidth parameter (5.34) and for readability issues, the optimisation problem (5.18) is recalled here:

$$\varepsilon_k^\star = \underset{\varepsilon_k \in \mathcal{D}}{\operatorname{argmin}} \left(\widehat{\mathrm{ESS}}(\varepsilon_k) - N_{\mathrm{th}} \right)^2 \tag{5.35}$$

For an iteration of the optimisation (i.e., for ε_k fixed), the computation load is described in Table 5.1.

Table 5.1: The flops complexity associated with the A2BC bandwidth parameter.

Instruction	Computation load
$\tilde{w}_{\varepsilon_k}$ (5.17b)	$N(c_{\mathcal{K}} + c_p + c_q + 3)$
Normalisation	$2N - 1$
$\widehat{\mathrm{ESS}}(\varepsilon_k)$	$2N$
$\left(\widehat{\mathrm{ESS}}(\varepsilon_k) - N_{\mathrm{th}}\right)^2$	$2N$

The load of the approximate effective sample size $\widehat{\mathrm{ESS}}(\varepsilon_k)$ includes the computation of the unnormalised weights $\tilde{w}$ (5.17b) and their normalisation.

When the proposal density is equal to the prior density, the computational load of the unnormalised weights becomes $c_{\mathcal{K}} + 1$. Thus, for bootstrap filters, the load of the bandwidth selection is:

$$c_\varepsilon = c_{\mathcal{D}} \left(N(c_{\mathcal{K}} + 7) - 1\right) \tag{5.36}$$

The number of calls of the objective function $c_{\mathcal{D}}$, and the number of particles N in the computational load of the bandwidth parameter ε, are the terms of greatest concern when looking at the load of the bandwidth parameter. However, since the cost is polynomial, the computation can be performed in real time on Graphics Processing Unit [12]. Moreover, the parallel processing of the bandwidth parameter optimisation (5.18) allows a significant decrease in computation time. The load of the bandwidth parameter can also be significantly reduced by verifying two values of the efficient sample size criterion. As discussed in Section 5.1.1, when $\mathrm{ESS}(\varepsilon_k^+) < N_{\mathrm{th}}$ and $\mathrm{ESS}(1) > N_{\mathrm{th}}$, the chosen bandwidth parameter $\varepsilon_k^\star$ is ε_k^+ and 1 respectively. The computational load is added to the total load of the algorithm only when $\mathrm{ESS}(1) < N_{\mathrm{th}} < \mathrm{ESS}(\varepsilon_k^+)$.

5.3 NUMERICAL STUDY: APPLICATION TO TERRAIN AIDED NAVIGATION

This section makes a comparison between classical Particle Filters (PFs) and their A2BC versions. In Section 5.3.1, the chosen dynamical model is the double integrator model (4.22), and the observation model is the Cartesian measurements for the multi-beam telemeter (4.27) and gravimeter (4.32).

The following points are to be checked:

1. The impact of the A2BC method on the accuracy (i.e., on the RMSE, see Section 4.4.1);
2. The impact of the A2BC method on the non-convergence rate (see Section 4.4.3).

It can be expected that A2BC filters would yield lower RMSE and non-convergence rate. The contribution of the gravimeter to the accuracy of the estimation and the robustness to uncertainties (measured by the non-convergence rate) is highlighted. The impact of the A2BC method is also studied in Section 5.3.2 for the inertial errors model. The chosen observation model is the geographical measurements for the multi-beam telemeter (4.26) and gravimeter (4.31).

To compare the conventional filters with their A2BC version, a Cauchy kernel with constant bandwidth parameter is chosen for RPF and RBPF likelihoods. The constant bandwidth is equal to the sensor error standard deviation. A2BC-RPF and A2BC-RBPF likelihoods are also Cauchy kernels but the bandwidth is adaptively determined according to the A2BC method previously described (see Section 5.1.1). The linear part of RBPF and A2BC-RBPF is the velocity on the three axis and the nonlinear part is the position on the three axis.

5.3.1 *Double Integrator Model*

The impact of the A2BC method is shown on a first dataset for an AUV equipped only with a multi-beam telemeter and an embedded map of the seabed. The joint contribution of the A2BC method and gravimetric sensor is shown for a second dataset.

The Contribution of the A2BC Method

The reference trajectory is simulated using the parameters presented in Table 5.2 and is shown in Figure 5.2. Figure 5.2 illustrates estimated trajectories given by a single Monte Carlo run of the two proposed algorithms (i.e., A2BC-RPF and A2BC-RBPF) for the same initial errors and measurements realisations. The estimated trajectories converge quickly towards the trajectory of reference. The simulation parameters are summarised in Table 5.2.

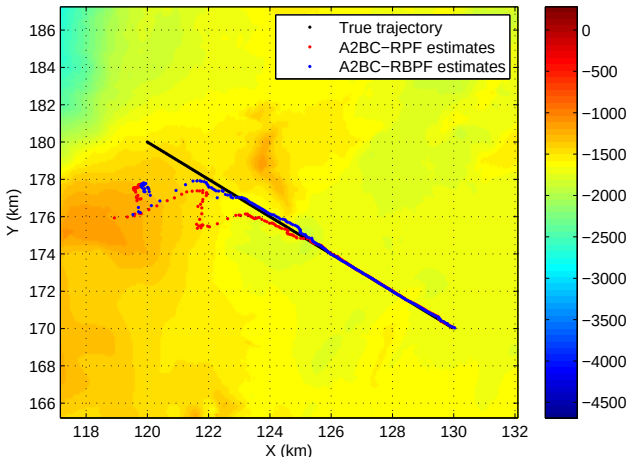

Figure 5.2: Comparison between the true trajectory and the trajectories estimated by a single Monte Carlo run of A2BC-RPF and A2BC-RBPF.

Table 5.2: Simulation settings.

State-space model parameters	Value
Number of Monte Carlo runs	100
Sampling period	$\Delta t = 10$ s
Number of bathymetric measurements	200
Trajectory duration	33 min
Number of beams	$m = 3$
Number of particles	$N = 5000$
Resampling threshold	$N_{th} = 0.75\ N$
Regularisation bandwidth parameter	$\mu = 0.3$
Initial position	$[120000, 180000, -100]^\top$ m
Initial velocity	$[5, 5, 0.05]^\top$ m s^{-1}
Initial uncertainty in position (st.d.)	$[1000, 1000, 100]$ m
Initial uncertainty in velocity (st.d.)	$[0.5, 0.5, 0.5]$ m s^{-1}
Process noise in position (st.d.)	$[3, 3, 0.3]$ m
Process noise in velocity (st.d.)	$[0.015, 0.015, 0.0015]$ m s^{-1}
Measurements	
Error of each beam range (st.d.)	$\sigma_{mb} = 10$ m
A2BC parameter	
Domain of the bandwidth parameter	$\mathcal{D} = [1; 3]$

Figures 5.3 and 5.4 show RMSEs obtained with the two A2BC filters (A2BC-RPF and A2BC-RBPF) and the two conventional filters (RPF and RBPF). Only convergent simulations are used to plot RMSEs on Figures 5.3 and 5.4. The non-convergence rates of the filters are provided in Table 5.3. The PCRB is displayed in Figures 5.3 and 5.4. The PCRB is approximated over 100 state samples at each time-step. For this reason, RMSEs can be lower than the PCRB at some point.

In Figure 5.3, RMSEs convergence occurs around 20 min for the horizontal position and around 15 min for the vertical position. RMSEs follow the tendency of the approximated PCRB. In Figure 5.3, the jumps in the horizontal position RMSE of A2BC-RPF between 10 min and 17 min are due to map ambiguities. The A2BC method leads the filter to wait for the ambiguity to be resolved before converging. The RMSE of A2BC-RPF still converges at the same time and with the same accuracy as the other filters.

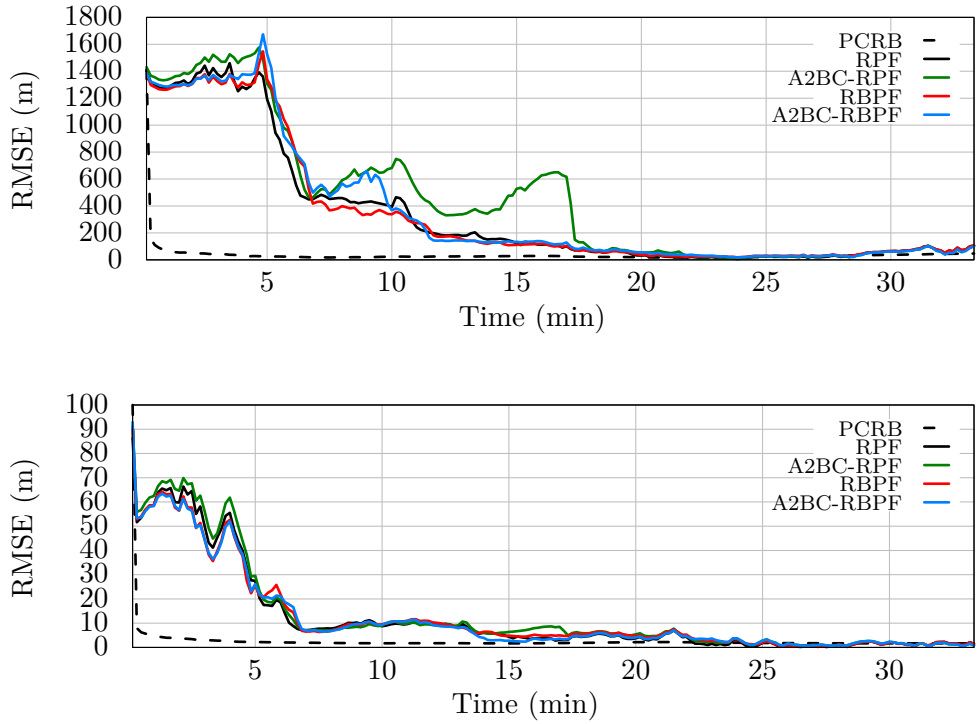

Figure 5.3: Plot of the RMSEs of the horizontal (upper plot) and vertical (lower plot) position in the bathymetric case for the double integrator model.

In Figure 5.4, RMSEs convergence occurs around 20 min for the horizontal velocity and around 15 min for the vertical velocity. RMSEs also follow the tendency of the approximated PCRB.

RMSEs at the end of the trajectory and non-convergence rates are provided in Table 5.3. RMSEs at the end of the trajectory are lower for the A2BC filters than for the conventional filters. This demonstrates that the boundedness of the asymptotic variance of the unnormalised weights (see Proposition 2) and of the local Monte Carlo error (see Proposition 3) has an impact on the estimation error in practice.

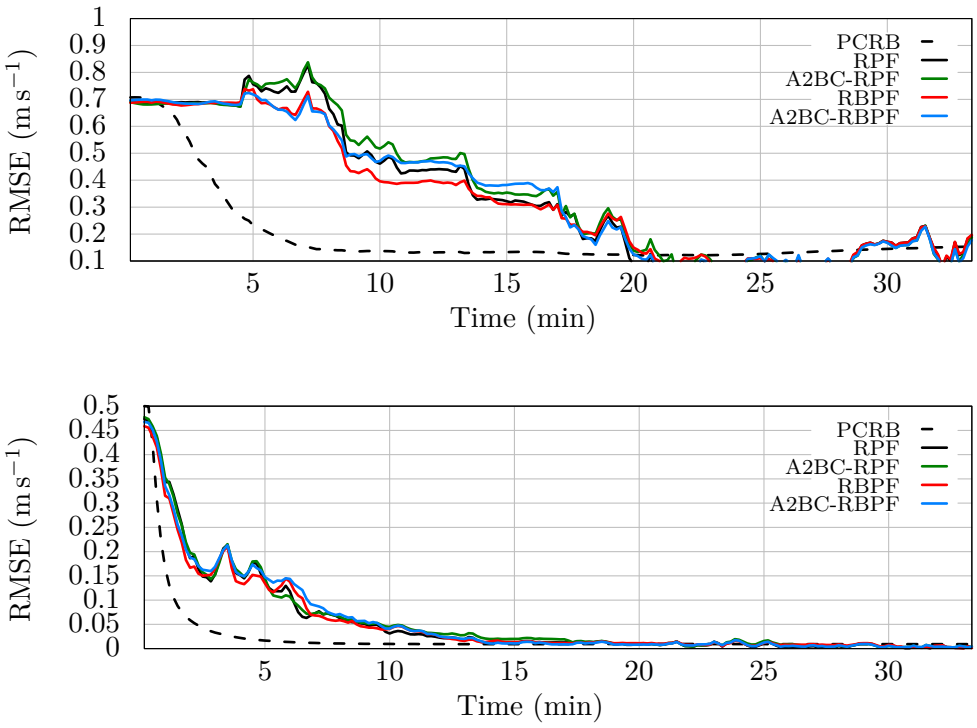

Figure 5.4: Plot of the RMSEs of the horizontal (upper plot) and vertical (lower plot) velocity in the bathymetric case for the double integrator model.

For 100 Monte Carlo simulations, the non-convergence rate decreases from 8 to 3 by adding the A2BC method in the RPF and from 10 for the RBPF to 6 for the A2BC-RBPF. This tends to show that A2BC filters are more robust to measurement ambiguities than their classic counterparts.

Table 5.3: Table of RMSEs at the end of the trajectory (in meters for the position and in meters per second for the velocity), and of the non-convergence rate.

		RPF	A2BC-RPF	RBPF	A2BC-RBPF
Last RMSE	p^h	106.53	104.67	105.23	99.90
	p^z	1.47	0.83	1.46	0.78
	v^h	0.795	0.188	0.197	0.177
Non-convergence rate (%)		8	**3**	10	**6**

The following section illustrates the joint contributions of the A2BC method and of the additional sensor: the atomic gravimeter.

The Gravimeter contribution

The reference trajectory is simulated using the parameters presented in Table 5.4 and is shown in Figure 4.9. The simulation parameters are summarised in Table 5.4.

Table 5.4: Simulation settings.

State-space model parameters	Value
Number of Monte Carlo runs	100
Sampling period	$\Delta t = 1$ s
Number of bathymetric measurements	600
Number of gravimetric measurements	40
Trajectory duration	10 min
Number of beams	$m = 5$
Number of particles	$N = 1000$
Resampling threshold	$N_{th} = 0.75\ N$
Regularisation bandwidth parameter	$\mu = 0.3$
Initial position	$[110000, 140000, -100]^\top$ m
Initial velocity	$[5, 5, 0.05]^\top$ m s^{-1}
Initial uncertainty in position (st.d.)	$[2000, 2000, 200]$ m
Initial uncertainty in velocity (st.d.)	$[0.5, 0.5, 0.05]$ m s^{-1}
Process noise in position (st.d.)	$[3, 3, 0.3]$ m
Process noise in velocity (st.d.)	$[0.015, 0.015, 0.0015]$ m s^{-1}
Measurements	
Error of each beam range (st.d.)	$\sigma_{mb} = 10$ m
Gravimeter error (st.d.)	$\sigma_{ga} = 0.3$ mGal
A2BC parameter	
Domain of the bandwidth parameter	$\mathcal{D} = [1; 2]$

Figures 5.5 and 5.6 show RMSEs obtained with the RPF and its A2BC version (A2BC-RPF). When the AUV is equipped with only the multi-beam telemeter, the term "Bathymetry" is added before the name of the filters. When the term "Fusion" is added before the name of the filters, the AUV is equipped with the multi-beam telemeter and the atomic gravimeter.

Only convergent Monte Carlo simulations are used to plot RMSEs on Figures 5.5 and 5.6. The non-convergence rates of the filters are provided in Table 5.8. The PCRB is displayed in Figure 5.5 and is approximated over 100 state samples at each time-step.

In Figure 5.5, RMSEs convergence occurs around 4 min for all filters. RMSEs follow the tendency of the approximated PCRB. At the beginning of the trajectory, RMSEs of the RPF and A2BC-RPF in the "Fusion" case are lower than the RMSEs of the same filters in the "Bathymetry" case which demonstrates the interest of fusing gravimetry based measurement to bathymetric information.

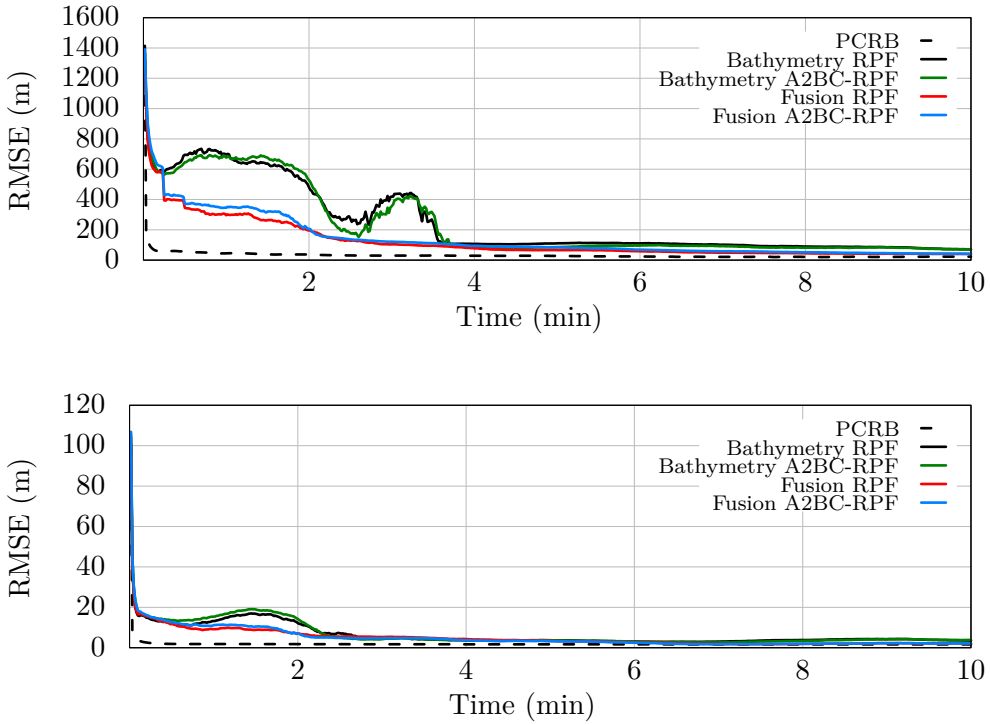

Figure 5.5: Plot of RMSEs of the horizontal (upper plot) and the vertical position (lower plot) for the double integrator model. "Bathymetric" corresponds to the bathymetric aided navigation and "Fusion" to the bathymetry and gravimetry fusion scenario.

In Figure 5.6, for both cases, RMSEs for the horizontal velocity converge slowly towards the value 0.015 $\mathrm{m\,s^{-1}}$. RMSEs for the vertical velocity converge at the end of the trajectory. The use of the atomic gravimeter, in addition to the multi-beam telemeter, allows a significant improvement of RMSEs for both filters. The gap is particularly large when considering the horizontal velocity. This significant difference between RMSEs is due to the explicit contribution of the x-axis velocity in the measurement equation of the atomic gravimeter (4.32).

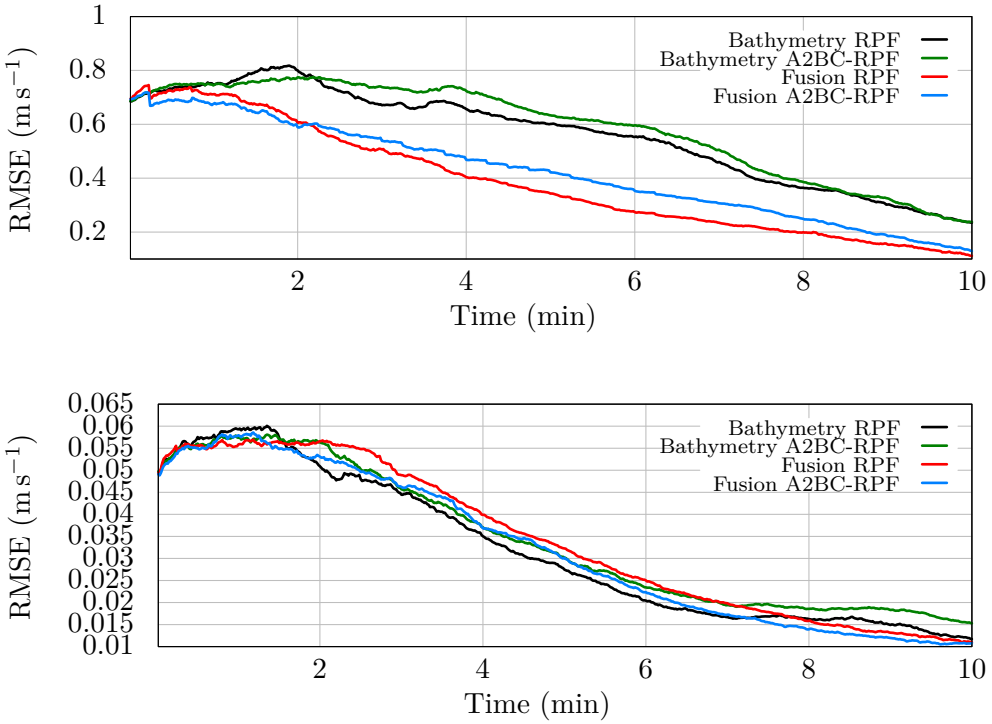

Figure 5.6: Plot of RMSEs of the horizontal (upper plot) and the vertical velocity (lower plot) for the double integrator model. "Bathymetric" corresponds to the bathymetric aided navigation and "Fusion" to the bathymetry and gravimetry fusion scenario.

RMSEs for the x-axis velocity are shown in Figure 5.7. The accurate estimation of the x-axis velocity leads to an overall improvement of the accuracy of the other state parameters estimates. Thus, the fusion strategy presented in Section 4.2.3 improves the quality of filter estimates.

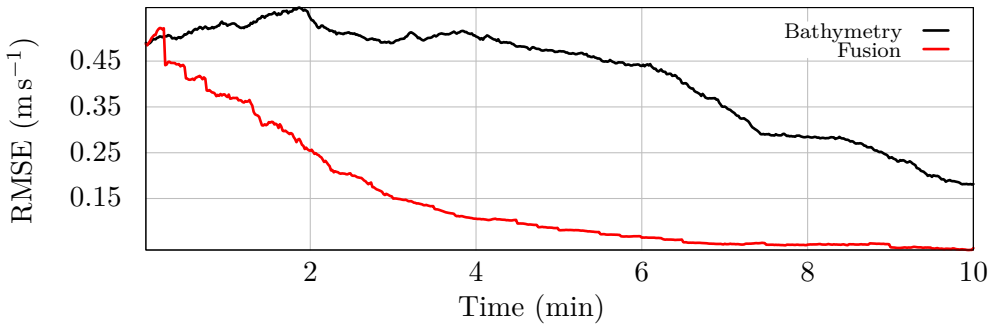

Figure 5.7: Plot of RPF RMSEs of the velocity on the x-axis. "Bathymetric" corresponds to the bathymetric aided navigation and "Fusion" to the bathymetry and gravimetry fusion scenario.

For 100 Monte Carlo simulations, the non-convergence rates of RPF and RBPF-RPF in the "Bathymetry" and "Fusion" cases are shown in Figure 5.8. Overall, the non-convergence rates are lower for the filters in the "Fusion" case. For the RPF, the non-convergence rate decreases from 12 in the "Bathymetry" case to 7 in the "Fusion" case. The fusion with the atomic gravimeter effectively reduces the non-convergence rate. The contribution of the A2BC method is also illustrated as smaller non-convergence rates are observed for the A2BC filters. In the "Bathymetry" case, the non-convergence rate decreases from 12 to 8 with the A2BC method. In the "Fusion" case, the non-convergence rate decreases from 7 for the RPF to 3 for the A2BC-RPF. This significant decrease in the non-convergence rates was expected as the A2BC method optimises the impact of the weights correction.

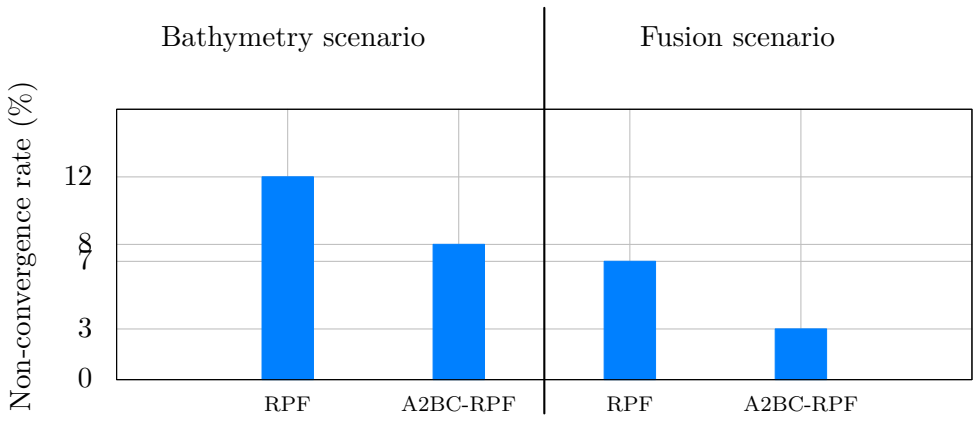

Figure 5.8: Non-convergence rates of the RPF and its A2BC version for 100 Monte Carlo simulations. "Bathymetry scenario" corresponds to the bathymetric aided navigation and "Fusion scenario" to sensor fusion case.

The multi-beam telemeter and atomic gravimeter fusion provides better estimates through the direct observation of the x-axis velocity in the gravimeter measurement equation.

The A2BC method brings more robustness to non-linearities and measurement ambiguities than classical filters as the non-convergence rates are significantly reduced. In the next section, A2BC filters are tested on a more realistic navigation model.

5.3.2 *IMU Error Drift Model*

The IMU error drift model is described in Chapter 4. The inertial errors dynamics are given by a linear model (see equation (4.14)). The impact of the A2BC method is shown for an AUV equipped with an IMU, a multi-beam telemeter and an embedded map of the seabed. The IMU measures the vehicle's acceleration and angular rate in the inertial frame, via respectively three accelerometers and three gyrometers. These measurements are integrated to retrieve the vehicle state. The accelerometers and gyrometers parameters of the IMU are summarised in Table 5.5, where bias, scale factors, and standard variations are the same for the three axis.

Table 5.5: IMU settings.

Gyrometers parameters	Value
Bias	$b^{\omega} = 0.005\ ^{\circ}\,\mathrm{h}^{-1}$
Correlation period	$\tau_a = 3000$ s
Scale factor	$K^{\omega} = 10^{-6}$
Random walk (st.d.)	$2.10^{-5}\ ^{\circ}/\sqrt{\mathrm{h}}$
Accelerometers parameters	**Value**
Bias	$b^{a} = 1.10^{-5}$ g
Correlation period	$\tau_{\omega} = 3000$ s
Scale factor	$K^{a} = 10^{-6}$
Random walk (st.d.)	$3.10^{-4}\ ^{\circ}/\mathrm{s}/\sqrt{\mathrm{h}}$

A high quality IMU is described in Table 5.5 as the bias of the gyrometers is small. An example of the IMU drift is shown in Figure 5.9. The drift is due to the sensors errors and the integration of the initial error.

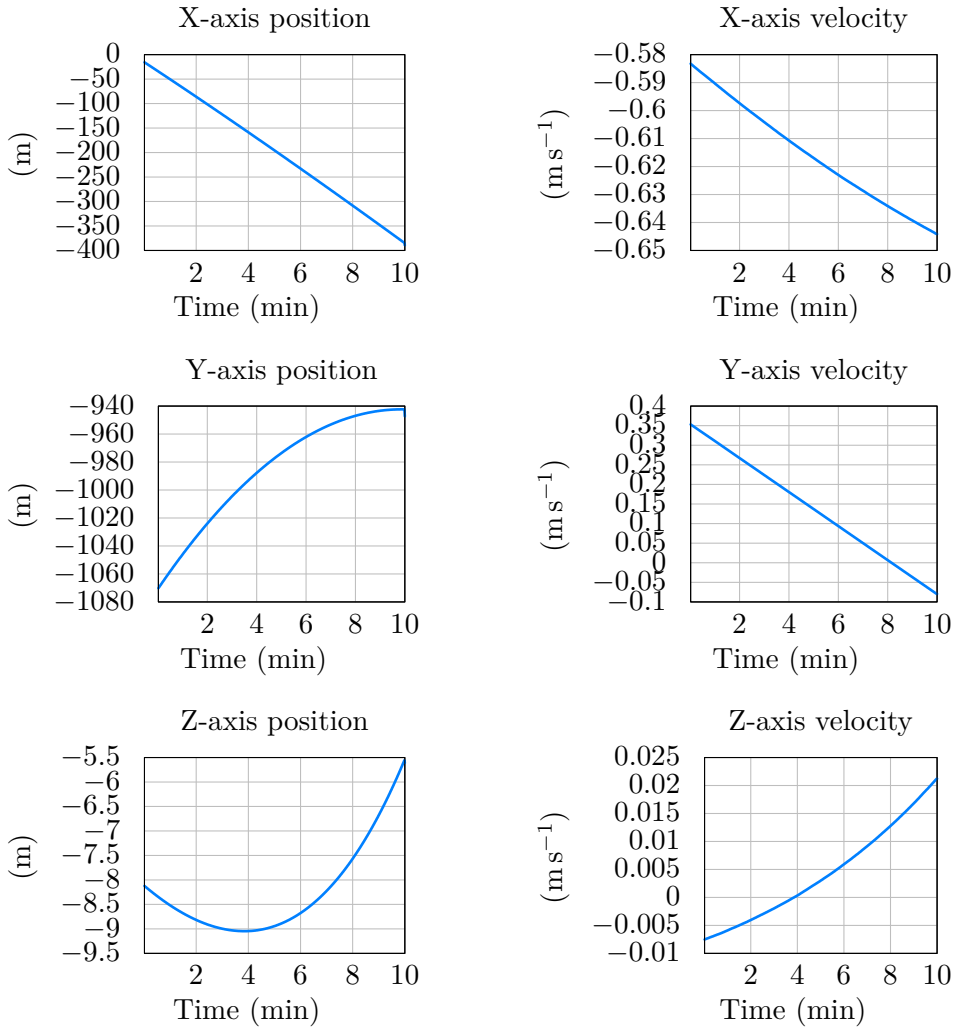

Figure 5.9: An example of IMU drift.

The reference trajectory is simulated using the parameters presented in Table 5.4 and is shown in Figure 4.9. The parameters of the filters are also given by Table 5.4. The state vector is composed of the position and the velocity.

Figures 5.10 and 5.11 shows RMSEs obtained with the RPF and the A2BC-RPF. Only convergent simulations are used to plot RMSEs on Figures 5.10 and 5.11. RMSEs decrease with time for both the horizontal and the vertical position and velocity. After 3 min of trajectory, RMSEs of the horizontal position converge around 100 m. RMSEs of the vertical position converge around 5 m before 3 min of trajectory. For the horizontal and the vertical position, the RPF RMSE converges faster than the A2BC-RPF RMSE.

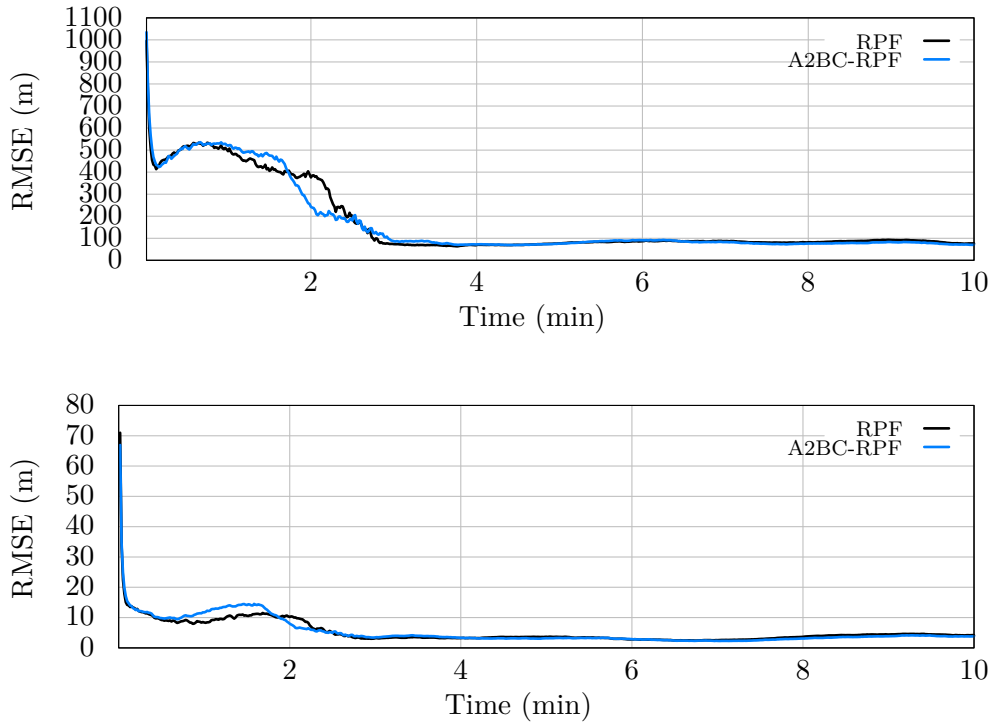

Figure 5.10: Plot of RMSEs for the horizontal (upper plot) and the vertical position (lower plot) for the IMU error drift model.

For the horizontal velocity, the convergence occurs at the end of the trajectory. RMSEs of the horizontal velocity reach 0.4 $\mathrm{m\,s^{-1}}$. RMSEs of the vertical velocity converge after 2 min of trajectory around the value 0.03 $\mathrm{m\,s^{-1}}$. The gaps between RMSEs of the RPF and RMSEs of the A2BC-RPF are small.

For 100 Monte Carlo simulations, the non-convergence rates are 16 for the RPF and 5 for the A2BC-RPF. We retrieve the same conclusions as in the previous simulations. The A2BC filter is more robust to non-linearities and measurement ambiguities as the non-convergence rate is significantly reduced.

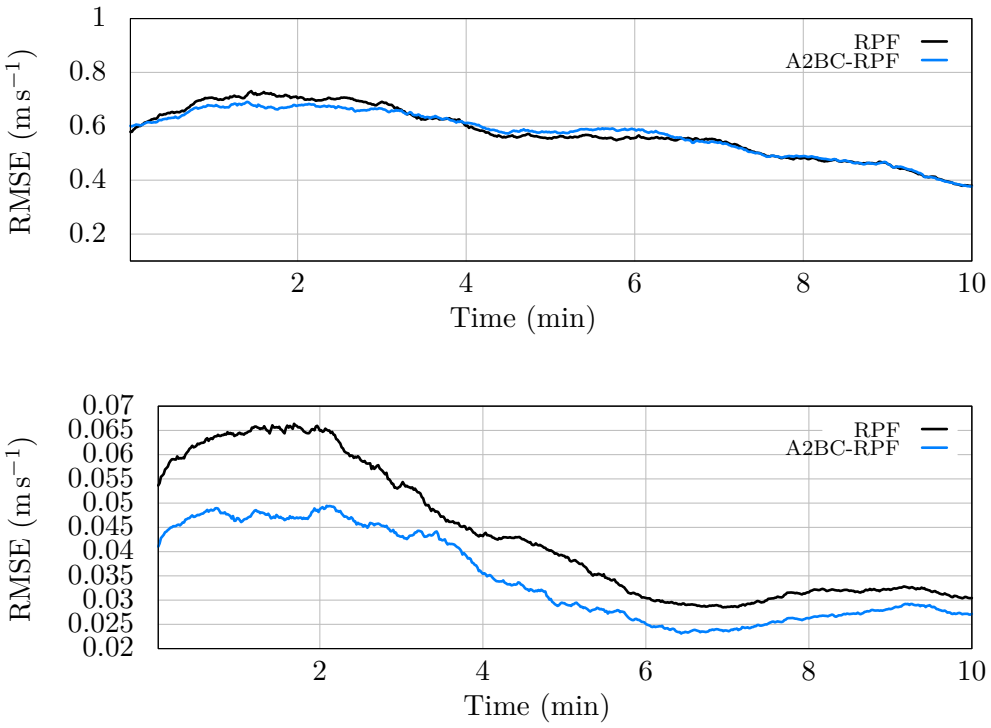

Figure 5.11: Plot of RMSEs for the horizontal (upper plot) and the vertical velocity (lower plot) for the IMU error drift model.

5.4 SUMMARY

This chapter introduced a method based on the ABC, called A2BC. The method A2BC, by adapting the support of the likelihood, ensures the consistency between the likelihood and the prior density. A2BC prevent the filter from diverging by optimising the weight correction impact. This theoretically allows the filter to tackle severe multi-modalities. The proposed method was applied in two conventional filters: RPF and RBPF.

A2BC filters were evaluated on the double integrator model and the inertial errors model. In practice, a significant decrease was shown in terms of non-convergence rate for both models. Simulation results demonstrate that the A2BC method significantly increases the robustness to non-linearities and measurement ambiguities, and the accuracy of the estimations, while remaining computationally efficient.

The contribution of the multi-beam telemeter and atomic gravimeter fusion was also demonstrated. The sensors fusion provides better estimates through the direct observation of the x-axis velocity in the gravimeter measurement equation and limits the position drift of the IMU.

6

INTERACTING WEIGHTED ENSEMBLE KALMAN FILTER

The choice of the proposal density is crucial as it ensures the correct functioning of particle algorithms. The proposal density should minimise the Monte Carlo variance. Unfortunately, the optimal proposal density depends on the target density (i.e., the conditional density). The advantage of the Weighted Ensemble Kalman Filter (WEnKF) [69] lies in the choice of the proposal density (see Section 3.3.4). WEnKF can be seen as a particle filter with a Gaussian proposal density calculated by a Kalman filter. However, in some cases, the support of the proposal density of the WEnKF does not overlap with the support of the conditional density, which causes a deterioration of the Monte Carlo estimate and may lead the filter to diverge.
To tackle this issue, it is possible to empirically enlarge the support of the proposal density so that it contains the support of the conditional density.

This chapter introduces several contributions to the WEnKF field:

- Section 6.2.1 determines a condition on the covariance of the proposal density so that it guarantees a finite asymptotic variance of the unnormalized weights.
- Section 6.2.2 derives an analytic formulation that satisfies this condition.

The resulting algorithm is called the Interacting Weighted Ensemble Kalman Filter (IWEnKF).

Section 6.3 concludes this chapter by comparing the performances in terms of RMSE and non-convergence rate of the IWEnKF with the conventional WEnKF on a TAN example.

6.1 PRINCIPLE

The rewriting of the Ensemble Kalman Filter (EnKF) correction equation allows us to include the WEnKF within the scope of PFs (see Section 3.3.4). Unlike bootstrap PFs, the WEnKF leads to a Gaussian proposal density (3.47), recalled here:

$$q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}^i, \boldsymbol{Y}_k) = \mathcal{N}(\boldsymbol{X}_k; \hat{\boldsymbol{\mu}}_k^i, \hat{\boldsymbol{P}}_k) \tag{6.1}$$

where

$$\hat{\boldsymbol{\mu}}_k^i = (\boldsymbol{I}_d - \boldsymbol{K}_k\boldsymbol{H}_k)\, b_k(\boldsymbol{X}_{k-1}^i) + \boldsymbol{K}_k\, \boldsymbol{Y}_k \tag{6.2a}$$

$$\hat{\boldsymbol{P}}_k = (\boldsymbol{I}_d - \boldsymbol{K}_k\boldsymbol{H}_k)\, \boldsymbol{Q}_k\, (\boldsymbol{I}_d - \boldsymbol{K}_k\boldsymbol{H}_k)^\top + \boldsymbol{K}_k\boldsymbol{R}_k\boldsymbol{K}_k \tag{6.2b}$$

The WEnKF proposal density is suboptimal as it does not minimises the variance of the weights (see Section 3.3.1). In some cases, the obtained proposal density is not close enough to the conditional density and thus the filter may diverge. To tackle this issue, it is possible to empirically enlarge the support of the proposal density so that it contains the support of the conditional density. As the proposal density is Gaussian, enlarging the support of the density is equivalent to enlarge its covariance. The idea behind the proposed algorithm is to replace the covariance $\hat{\boldsymbol{P}}_k$ to ensure that the asymptotic variance of the unnormalised weights remains bounded.

Thereafter, the prior density $p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})$ is considered Gaussian:

$$p(\boldsymbol{X}_k|\mathcal{Y}_{k-1}) \propto \exp\left(-\frac{1}{2}\,(\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k)^\top \tilde{\boldsymbol{P}}_k^{-1}(\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k)\right) \tag{6.3}$$

where the mean and covariance of the prior density are respectively $\tilde{\boldsymbol{X}}_k$ and $\tilde{\boldsymbol{P}}_k$.

6.2 THEORETICAL DESCRIPTION

In this section, we first determine a condition on the covariance of the proposal density so that it guarantees that the asymptotic variance of the unnormalized weights is finite (Section 6.2.1). A small asymptotic variance of the unnormalized weights guarantees a more accurate Monte Carlo estimate. We then derive an analytic formulation of the IWEnKF that satisfies this condition (Section 6.2.2).

6.2.1 *Upper Bound of the Asymptotic Variance of the Unnormalized Weights*

Under Gaussian assumptions, Proposition 4 states that the asymptotic variance of the unnormalised weights admits a finite upper bound under a condition on the covariance $\hat{\boldsymbol{P}}_k$ of the proposal density.

Proposition 4. *The prior density $p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})$ and the proposal density $q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)$ are assumed to be Gaussian with means $\tilde{\boldsymbol{X}}_k$ and $\hat{\boldsymbol{X}}_k$, and covariances $\tilde{\boldsymbol{P}}_k$ and $\hat{\boldsymbol{P}}_k$ respectively. The likelihood is assumed to be bounded and the product between the likelihood and prior density is assumed to be integrable. If $\hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k$ is positive definite then the asymptotic variance of the unnormalized weights is finite:*

$$\hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k > 0 \Leftrightarrow \mathbb{V} < +\infty \tag{6.4}$$

Proof. The asymptotic variance of the unnormalized weights (3.56) is given by:

$$\mathbb{V} = \frac{\int \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k)^2\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})^2}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)}\, d\boldsymbol{X}_k}{\left(\int p(\boldsymbol{Y}_k|\boldsymbol{X}_k)\, p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})\, d\boldsymbol{X}_k\right)^2} - 1 \tag{6.5}$$

We assume that the likelihood $p(\boldsymbol{Y}_k|\boldsymbol{X}_k)$ is bounded, and the product between the likelihood and prior density is integrable. The asymptotic variance $\mathbb{V}$ is bounded if the numerator is bounded:

$$\int \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k)^2 p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})^2}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \, d\boldsymbol{X}_k < +\infty \tag{6.6}$$

By taking the supremum, it comes:

$$\int \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k)^2 p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})^2}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \, d\boldsymbol{X}_k \leq \sup_{\boldsymbol{X}_k \in \mathbb{R}^d} p(\boldsymbol{Y}_k|\boldsymbol{X}_k) \sup_{\boldsymbol{X}_k \in \mathbb{R}^d} \left(\frac{p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \right) \int p(\boldsymbol{Y}_k|\boldsymbol{X}_k) p(\boldsymbol{X}_k|\mathcal{Y}_{k-1}) \, d\boldsymbol{X}_k \tag{6.7}$$

Then, a sufficient condition for the asymptotic variance of the unnormalized weights to be bounded is:

$$\sup_{\boldsymbol{X}_k \in \mathbb{R}^d} \left(\frac{p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \right) < +\infty \tag{6.8}$$

The ratio between the proposal and prior densities is bounded if and only if the logarithm of the ratio is bounded:

$$\log \left(\frac{p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \right) < +\infty \tag{6.9}$$

By replacing the densities $p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})$ and $q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)$ by their expressions (see equations (6.3) and (6.1)), it follows:

$$\log \left(\frac{p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \right) \propto \frac{1}{2} (\boldsymbol{X}_k - \hat{\boldsymbol{X}}_k)^\top \hat{\boldsymbol{P}}_k^{-1} (\boldsymbol{X}_k - \hat{\boldsymbol{X}}_k) - \frac{1}{2} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k)^\top \tilde{\boldsymbol{P}}_k^{-1} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k) \tag{6.10}$$

By simultaneously adding and subtracting $\tilde{\boldsymbol{X}}_k$ and pooling the terms, the logarithm (6.10) can be written as follows:

$$\log \left(\frac{p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}, \boldsymbol{Y}_k)} \right) \propto \frac{1}{2} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k + \tilde{\boldsymbol{X}}_k - \hat{\boldsymbol{X}}_k)^\top \hat{\boldsymbol{P}}_k^{-1} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k + \tilde{\boldsymbol{X}}_k - \hat{\boldsymbol{X}}_k) - \frac{1}{2} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k)^\top \tilde{\boldsymbol{P}}_k^{-1} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k) \tag{6.11}$$

$$\propto \frac{1}{2} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k)^\top \hat{\boldsymbol{P}}_k^{-1} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k) + \frac{1}{2} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k)^\top \hat{\boldsymbol{P}}_k^{-1} (\tilde{\boldsymbol{X}}_k - \hat{\boldsymbol{X}}_k) + \frac{1}{2} (\tilde{\boldsymbol{X}}_k - \hat{\boldsymbol{X}}_k)^\top \hat{\boldsymbol{P}}_k^{-1} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k) - \frac{1}{2} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k)^\top \tilde{\boldsymbol{P}}_k^{-1} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k) \tag{6.12}$$

$$\propto -\frac{1}{2} (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k)^\top \left(\tilde{\boldsymbol{P}}_k^{-1} - \hat{\boldsymbol{P}}_k^{-1} \right) (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k) + \boldsymbol{V}_k^\top (\boldsymbol{X}_k - \tilde{\boldsymbol{X}}_k) \tag{6.13}$$

where $\boldsymbol{V}_k = \hat{\boldsymbol{P}}_k^{-1}(\tilde{\boldsymbol{X}}_k - \hat{\boldsymbol{X}}_k)$.

By taking $||\boldsymbol{X}_k|| \to +\infty$, it comes:

$$\lim_{||\boldsymbol{X}_k||\to+\infty} \log\left(\frac{p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1},\boldsymbol{Y}_k)}\right) = -\infty \quad \text{if} \quad \tilde{\boldsymbol{P}}_k^{-1} - \hat{\boldsymbol{P}}_k^{-1} > 0 \tag{6.14}$$

As the logarithm of the ratio between the proposal and prior densities tends towards $-\infty$ when $||\boldsymbol{X}_k||$ tends to $+\infty$, the ratio tends towards 0 and is thus bounded.

If the condition $\hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k > 0$ is satisfied, the supremum is bounded:

$$\sup_{\boldsymbol{X}_k\in\mathbb{R}^d}\left(\frac{p(\boldsymbol{X}_k|\mathcal{Y}_{k-1})}{q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1},\boldsymbol{Y}_k)}\right) < +\infty \tag{6.15}$$

Therefore, the asymptotic variance of the unnormalised weight is bounded. □

In order to satisfy condition (6.4) in the Proposition 4, the covariance of the proposal density $\hat{\boldsymbol{P}}_k$ must be appropriately chosen. In the following section, we introduce a method to determine the covariance of the proposal density in the WEnKF framework which leads to the IWEnKF formulation.

6.2.2 *Interacting Weighted Ensemble Kalman Filter*

In practice, the proposal density given by the WEnKF is the Gaussian given by the equation (6.1). The covariance $\hat{\boldsymbol{P}}_k$ of the proposal density is not guaranteed to satisfy the conditions of Proposition 4: $(\hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k > 0)$. The initial covariance $\hat{\boldsymbol{P}}_k$ can be replaced with a new covariance $\hat{\boldsymbol{P}}_k^\star$ so that $\hat{\boldsymbol{P}}_k^\star - \tilde{\boldsymbol{P}}_k > 0$. There is an infinite number of matrices $\hat{\boldsymbol{P}}_k^\star$ such that the condition $\hat{\boldsymbol{P}}_k^\star - \tilde{\boldsymbol{P}}_k > 0$ is respected. The idea behind the IWEnKF is to take the one that is as close as possible (in terms of the Frobenius norm) to the initial covariance $\hat{P}_k$ given by the WEnKF. This approach allows a greater overlap between the proposal and conditional densities, as illustrated in Figure 6.1, without loosing the information contained in the initial covariance.

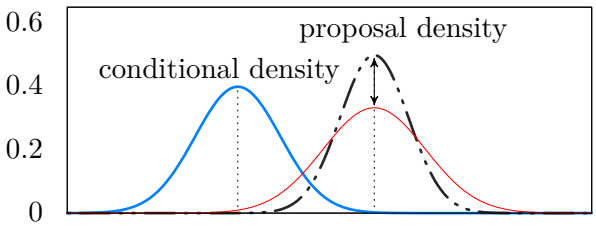

Figure 6.1: Support recovery between the conditional density and two proposal densities: the original WEnKF proposal density (dashed black line), and the IWEnKF proposal density (thin red line).

Proposition 5. *A covariance $\hat{\boldsymbol{P}}_k^\star$ of the proposal density that satisfies the condition of Proposition 4, i.e., $\hat{\boldsymbol{P}}_k^\star - \tilde{\boldsymbol{P}}_k > 0$, is:*

$$\hat{\boldsymbol{P}}_k^\star = \frac{1}{2}(\hat{\boldsymbol{P}}_k + \tilde{\boldsymbol{P}}_k + \boldsymbol{H}_k) \tag{6.16}$$

where:

- *$\tilde{\boldsymbol{P}}_k$ (6.3) and $\hat{\boldsymbol{P}}_k$ (6.2b) are the covariance of respectively the prior and the WEnKF proposal densities.*
- *$\boldsymbol{H}_k$ is obtained from the polar decomposition of $\hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k$, i.e., $\hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k = \boldsymbol{U}_k\boldsymbol{H}_k$ with $\boldsymbol{U}_k^\top \boldsymbol{U}_k = \mathbf{I}_d$.*

Proof. Finding a covariance $\boldsymbol{P}_k$ such that the condition $\boldsymbol{P}_k - \tilde{\boldsymbol{P}}_k > 0$ is satisfied and $\boldsymbol{P}_k$ close to the initial covariance $\hat{\boldsymbol{P}}_k$ given by the WEnKF is equivalent to finding a symmetrical matrix $\boldsymbol{P}_k$ solution of the following constrained optimisation problem:

$$\min_{\boldsymbol{P}_k - \tilde{\boldsymbol{P}}_k > 0} ||\boldsymbol{P}_k - \hat{\boldsymbol{P}}_k||_F^2 \tag{6.17}$$

where $||.||_F$ is the Frobenius norm. In order to keep the information contained in the covariance of the proposal density fo the WEnKF (6.1), we want to find $\boldsymbol{P}_k$ close to $\hat{\boldsymbol{P}}_k$.

By simultaneously adding and subtracting $\tilde{\boldsymbol{P}}_k$ in the previous minimisation problem (6.17), it comes:

$$\min_{\boldsymbol{P}_k - \tilde{\boldsymbol{P}}_k > 0} ||\boldsymbol{P}_k - \tilde{\boldsymbol{P}}_k - (\hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k)||_F^2 \tag{6.18}$$

By taking $\boldsymbol{V}_k = \boldsymbol{P}_k - \tilde{\boldsymbol{P}}_k$ and $\boldsymbol{A}_k = \hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k$, the optimisation problem (6.18) becomes:

$$\min_{\boldsymbol{V}_k > 0} ||\boldsymbol{V}_k - \boldsymbol{A}_k||_F^2 \tag{6.19}$$

It can be approximated by the nearest symmetric positive definite matrix, in terms of the Frobenius norm, by using Higham's theorem [39]:

$$\boldsymbol{V}_k^F = \frac{1}{2}(\boldsymbol{A}_k + \boldsymbol{H}_k) \tag{6.20}$$

where $\boldsymbol{H}_k$ is obtained from the polar decomposition of $\hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k$, i.e. $\hat{\boldsymbol{P}}_k - \tilde{\boldsymbol{P}}_k = \boldsymbol{U}_k\boldsymbol{H}_k$ with $\boldsymbol{U}_k^\top \boldsymbol{U}_k = \mathbf{I}_d$.

Thus, the solution of the optimisation problem (6.17) is:

$$\hat{\boldsymbol{P}}_k^\star = \boldsymbol{V}_k^F + \tilde{\boldsymbol{P}}_k = \frac{1}{2}(\hat{\boldsymbol{P}}_k + \tilde{\boldsymbol{P}}_k + \boldsymbol{H}_k) \tag{6.21}$$

□

Proposition 5 introduces a new choice of covariance for the proposal density that satisfies the conditions of Proposition 4. The IWEnKF uses this new covariance to compute the proposal density:

$$q(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}^i, \boldsymbol{Y}_k) = \mathcal{N}(\boldsymbol{X}_k; \hat{\boldsymbol{\mu}}_k^i, \hat{\boldsymbol{P}}_k^\star) \tag{6.22}$$

where $\hat{\boldsymbol{\mu}}_k^i$ is given by equation (6.2a) and $\hat{\boldsymbol{P}}_k^\star$ by (6.16). The IWEnKF guarantees a finite asymptotic variance of the unnormalized weights. The new covariance of the proposal density is obtained via a polar decomposition which can be polynomially computed [38]. In practice, this is likely to bring more robustness to nonlinearities and multi-modality. The algorithm of IWEnKF is described in Algorithm 6.1.

Algorithm 6.1 Interacting Weighted Ensemble Kalman Filter

Initialisation: The initial particles set is drawn as $\{\boldsymbol{X}_0^i\}_{i\in[1,N]}$ using the initial density $p(\boldsymbol{X}_0) \sim \mathcal{N}(\boldsymbol{X}_0; 0_d, \boldsymbol{P}_0)$ and the initial weights set $\{w_0^i\}_{i\in[1,N]}$ is taken equal to $\frac{1}{N}$.
for each time-step k **do**
  Particle prediction: Compute the mean $\hat{\boldsymbol{\mu}}_k^i$ (6.2a) and covariance $\hat{\boldsymbol{P}}_k^*$ (6.16) of the proposal density. Draw the particles from the proposal density $\boldsymbol{X}_k^i \sim \mathcal{N}(\boldsymbol{X}_k; \hat{\boldsymbol{\mu}}_k^i, \hat{\boldsymbol{P}}_k^\star)$ (6.22).
  Particle correction: Update the weights

$$w_k^i \propto w_{k-1}^i \frac{p(\boldsymbol{Y}_k|\boldsymbol{X}_k^i)\; p(\boldsymbol{X}_k^i|\boldsymbol{X}_{k-1}^i)}{\mathcal{N}(\boldsymbol{X}_k^i; \hat{\boldsymbol{\mu}}_k^i, \hat{\boldsymbol{P}}_k^\star)}$$

  if a resampling criterion is satisfied, e.g., $\widehat{\mathrm{ESS}} < N_{\mathrm{th}}$, see (3.19) **then**
    Draw a new set of particles $\{\boldsymbol{X}_k^i\}_{i\in[1,N]}$ using a resampling method, e.g., the multinomial resampling (Algorithm 3.2).
    Reset the weights to $\frac{1}{N}$.
  end if
  Compute the linear state estimate $\hat{\boldsymbol{X}}_k$ (3.16) and its covariance $\boldsymbol{P}_k$ (3.17).
end for
Return the state estimate $\hat{\boldsymbol{X}}_k$ and its covariance $\boldsymbol{P}_k$, $\forall k$.

6.3 NUMERICAL RESULTS

To illustrate the behavior of the IWEnKF, it will be compared with the standard WEnKF on an underwater TAN example. The chosen dynamical model is the double integrator model 4.22 and the observation model is the Cartesian measurements equation of the multi-beam telemeter (4.27). Comparisons are done with the metrics described in Section 4.4. The simulation parameters are summarised in Table 6.1.

Table 6.1: Simulation settings.

State-space model parameters	Value
Number of Monte Carlo runs	50
Sampling period	$\Delta t = 5$ s
Number of bathymetric measurements	420
Trajectory duration	35 min
Number of beams	$m = 5$
Number of particles	$N = 3000$
Resampling threshold	$N_{th} = 0.75\ N$
Initial position	$[110000, 140000, -100]^\top$ m
Initial velocity	$[5, 5, 0.05]^\top$ m s^{-1}
Initial uncertainty in position (st.d.)	$[1000, 1000, 100]^\top$ m
Initial uncertainty in velocity (st.d.)	$[0.5, 0.5, 0.5]^\top$ m s^{-1}
Process noise in position (st.d.)	$[3, 3, 0.3]$ m
Process noise in velocity (st.d.)	$[0.02, 0.02, 0.002]$ m s^{-1}
Error of each beam range (st.d.)	$\sigma_{mb} = 10$ m

The PCRB is approximated over 300 state samples at each time-step. For this reason, RMSEs can be lower than the PCRB at some point. For one Monte Carlo simulation, Figure 6.2 illustrates the rapid convergence of a IWEnKF trajectory towards the true trajectory.

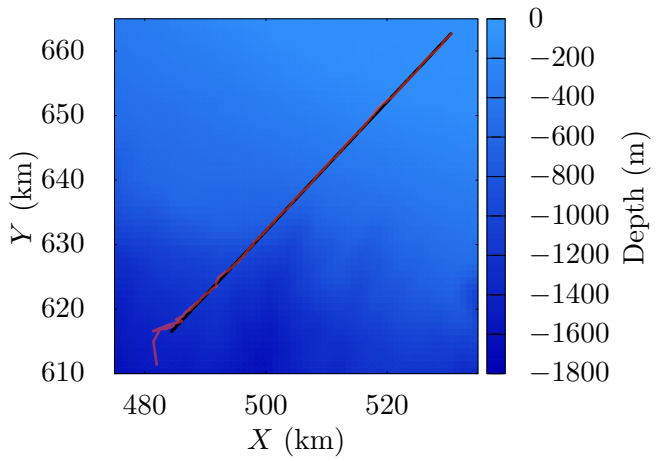

Figure 6.2: Bathymetric map of the Californian coast with the true trajectory (straight black line) and the trajectory estimated by a single Monte Carlo run of IWEnKF (red line).

Only convergent Monte Carlo simulations are used to plot RMSEs on Figure 6.3. The non-convergence rates of the algorithms are provided in Figure 6.4.

In Figure 6.3, the curves follow the tendency of the PCRB. Position and velocity RMSEs decrease with time and converge to a value close to the PCRB approximation. Although the accuracy of the filter estimates is similar after 20 min of trajectory, the IWEnKF converges faster than the WEnKF. Before the convergence of the filters (from 90 s, when the IWEnKF curve falls below the WEnKF curve, until 15 min), the average distance between the two curves is about 50 m for the horizontal position. For 50 Monte Carlo simulations, the non-convergence rates of both algorithms is shown in Figure 6.4. The non-convergence rate of the IWEnKF is twice as small as the non-convergence rate of the WEnKF. This significant decrease in the number of non-convergences was expected as the IWEnKF method guarantees a bounded asymptotic variance of the unnormalised weights and thus prevents weight degeneracy.

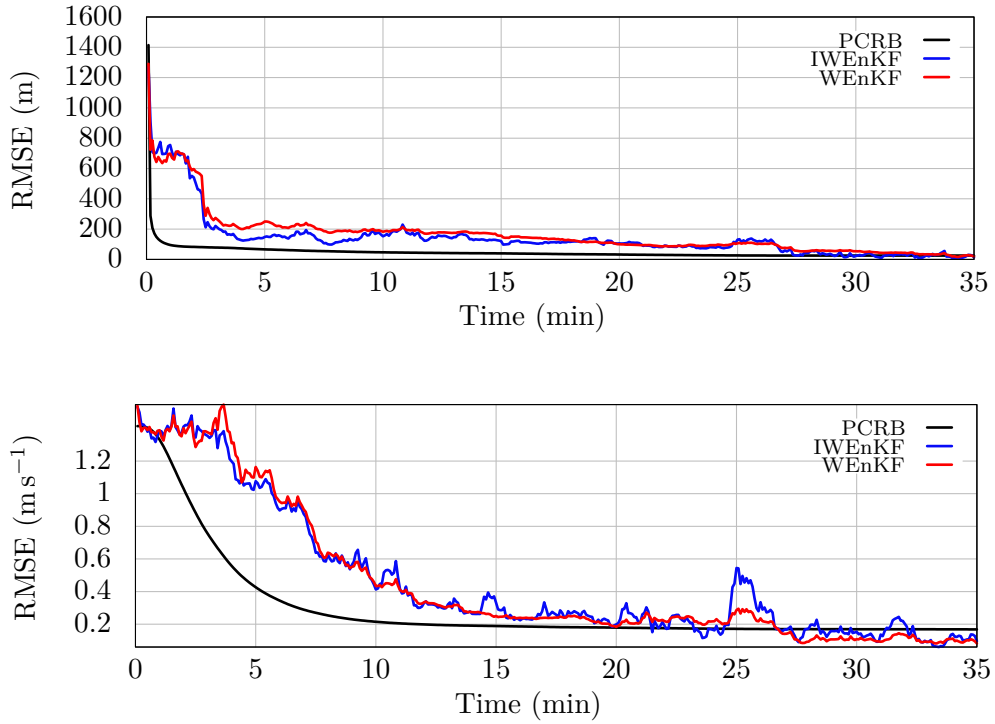

Figure 6.3: Plot of the PCRB and RMSEs for the horizontal position (upper plot) and the horizontal velocity (lower plot).

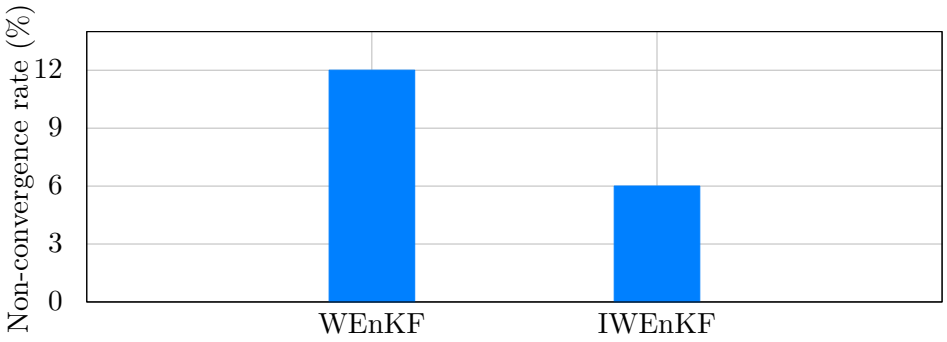

Figure 6.4: Histogram of the non-convergence rates of WEnKF and IWEnKF for 50 Monte Carlo simulations.

6.4 SUMMARY

This chapter introduced a weighted ensemble Kalman filter called IWEnKF. The proposed approach guarantees that the asymptotic variance of the unnormalised weights remains bounded which brings more robustness to nonlinearities in practice. IWEnKF is able to tackle non-linear dynamical models, as in the TAN problem. The guarantee of convergence of the IWEnKF is of strong interest for embedded applications on autonomous systems such as AUVs. Finally, the IWEnKF appears to be a good trade-off between robustness to measurements uncertainties and computational load as the polar decomposition can be polynomially computed. The IWEnKF was evaluated for the double integrator model and appears suitable for future integration in the inertial errors model with a multi-beam telemeter and a gravimeter. The study of the convergence of the IWEnKF in such a scheme will be tackled in future works.

7

CONCLUSION

In the case of non-linear measurements with severe ambiguities, the Kalman Filter (KF) and its derivatives fail to resolve the estimation problem. In many applications with strong nonlinearities, as the TAN scheme, the state estimation is resolved by PFs. The PFs successive Monte Carlo approximations may lead to the filter divergence.

To reduce the variance of the Monte Carlo approximation, two approaches were proposed.

- The first approach guaranteed the consistency between the likelihood and the prior density. We proposed a new family of filters, called A2BC filters. The A2BC method adapts the support of the likelihood in order to ensure a consistency between the likelihood and the prior density. The likelihood is approximated by a kernel density function that spreads out to ensure the consistency without flattening. The information delivered by the measurements is then preserved. The A2BC method reduces the occurrence of the resampling step and thus limits filter divergence.

- The second approach introduced a proposal density for which the support is assured to overlap the support of the conditional density. We introduced a state estimation algorithm, called IWEnKF. The IWEnKF proposed a proposal density for which the support overlap the support of the conditional density. This choice of proposal density minimises the Monte Carlo variance and thus provides better accuracy and robustness to nonlinearities.

The application investigated in this thesis was the underwater navigation of an AUV. AUV inertial navigation provides a solution of navigation that drifts with time due to imperfect sensors measurements. To correct the navigation drift, the IMU can be combined with external sensors. TAN provides a drift-free navigation tool. This application involves severe measurements ambiguities leading to the failure of many existing estimation methods. The presence of measurement ambiguities (e.g., to one terrain sensing measurement may correspond to several geographical areas of similar terrain profiles) increases the complexity of the state estimation.

The complexity of the state estimation was also increased by the introduction of an additional sensor: the atomic gravimeter. The addition of the atomic gravimeter has improved the estimates of the AUV state (especially for the velocity) and increased the robustness to nonlinearities and measurement ambiguities (the non-convergence rate decrease from 12% to 7% for the RPF).

The performance of A2BC filters (A2BC-RPF and A2BC-RBPF) was studied and compared to other PFs (RPF and RBPF) on an underwater TAN problem. The A2BC filters were tested on a realistic inertial model. The A2BC method has provided an increased accuracy of the estimation and more robustness to nonlinearities as the non-convergence rate is significantly reduces (the non-convergence rate decrease from 8% for the RPF to 3% for the A2BC-RPF).

The performance of the IWEnKF was studied and compared to the WEnKF on an underwater TAN problem. The IWEnKF has provided an increased accuracy of the estimation and more robustness to nonlinearities as the non-convergence rate is significantly reduces (the non-convergence rate decrease from 12% for the WEnKF to 6% for the IWEnKF).

Several future research directions can be drawn from this work.

- Compare the A2BC filters and IWEnKF with sensors measurement from test campaigns.
- It would be interesting to integrate the A2BC approach to the box particle filter [33, 59]. The result will allow a significant reduction of the computational cost compared to previous approaches.
- It would be interesting to use other Ensemble Kalman filters, e.g., the filter proposed by Sakov and Oke [77], in the IWEnKF to improve estimation accuracy.

a

BACKWARD NONLINEAR SMOOTHING DIFFUSIONS

Backward Nonlinear Smoothing Diffusions

Brian D.O. Anderson*[1], Adrian N. Bishop[2], Pierre Del Moral†[3], and Camille Palmier[4]

[1]Research School of Electrical, Energy and Material Engineering, Australian National University, Canberra, Australia; and also with Hangzhou Dianzi University, China, and Data61-CSIRO in Canberra, Australia.
[2]CSIRO & University of Technology Sydney (UTS), Australia
[3]INRIA, Bordeaux Research Center, Talence & CMAP, Polytechnique Palaiseau, France
[4]Institut de Mathématiques de Bordeaux (IMB), Bordeaux University & ONERA Palaiseau, France

Abstract

We present a backward diffusion flow (i.e. a backward-in-time stochastic differential equation) whose marginal distribution at any (earlier) time is equal to the smoothing distribution when the terminal state (at a latter time) is distributed according to the filtering distribution. This is a novel interpretation of the smoothing solution in terms of a nonlinear diffusion (stochastic) flow. This solution contrasts with, and complements, the (backward) deterministic flow of probability distributions (viz. a type of Kushner smoothing equation) studied in a number of prior works. A number of corollaries of our main result are given including a derivation of the time-reversal of a stochastic differential equation, and an immediate derivation of the classical Rauch-Tung-Striebel smoothing equations in the linear setting.

Keywords: Nonlinear filtering and smoothing; Kalman-Bucy filter; Rauch-Tung-Striebel smoother; particle filtering and smoothing; diffusion equations; stochastic semigroups; backward stochastic integration; backward Itô-Ventzell formula; time-reversed stochastic differential equations; Zakai and Kushner-Stratonovich equations.

Mathematics Subject Classification: 60G35; 62M20; 93E11; 93E14; 60J60.

1 Introduction

Let $(W_t, V_t) \in (\mathbb{R}^p \times \mathbb{R}^q)$ be a $(p+q)$-dimensional Brownian motion for finite $p, q \geq 1$. Consider a signal-observation model $(X_t, Y_t) \in (\mathbb{R}^m \times \mathbb{R}^n)$ given by the Itô stochastic differential equation:

$$\left\{ \begin{array}{rcl} dX_t & = & a_t(X_t)\, dt + \sigma_t(X_t)\, dW_t \\ dY_t & = & b_t(X_t)\, dt + \varsigma_t\, dV_t \end{array} \right. \tag{1.1}$$

for some measurable functions $\varsigma_t, a_t(x), \sigma_t(x), b_t(x)$ with appropriate dimensions. We set $Y_0 = 0$ and let X_0 be an initial random variable with absolute moments of any order. We let $\alpha_t(x) := \sigma_t(x)\sigma_t'(x)$, and $\beta_t := \varsigma_t\varsigma_t'$, where A' denotes the transpose of some matrix A.

*B.D.O. Anderson was supported by the Australian Research Council (ARC) via grant DP160104500 and grant DP190100887; and by Data61-CSIRO.

†P. Del Moral was supported in part by the Chair Stress Test, RISK Management and Financial Steering, led by the French Ecole Polytechnique and its Foundation and sponsored by BNP Paribas.

To avoid unnecessary technical details, we assume $\beta_t \geq \epsilon I$, for some $\epsilon > 0$, where I denotes the identity matrix. We also assume the drift and sensor functions $(a_u(x), b_u(x))$, as well as the diffusion matrix $\sigma_u(x)$, are smooth w.r.t. (u, x) and they have uniformly bounded derivatives w.r.t x of all order on $(u, x) \in [s, t] \times \mathbb{R}^m$, for any $s \leq t$.

These technical conditions ensure that the above stochastic differential equation (1.1) has a global solution (X_t, Y_t) in the sense of Itô. In addition, (X_t, Y_t) as well as the sensor function $b_t(X_t)$ have absolute moments of any order. The stochastic flow associated with the signal is also smooth w.r.t. its initial condition, and its derivatives have absolute moments of any order.

The filtering problem then consists of computing the conditional distribution π_t of the random signal states X_t of the signal given the sigma-field $\mathcal{Y}_t = \sigma(Y_s,\ s \leq t)$ generated by the observations. The smoothing problem is to compute the conditional distribution $\pi_{t,s}$ of the random signal states X_s given $\mathcal{Y}_t$, with $t \geq s$. With this notation, we have $\pi_{t,t} = \pi_t$.

The filtering and smoothing problems have been studied extensively, and the literature on this topic is too broad to survey in detail here; and a review of this type is beyond the rather narrow scope of our contribution. We may point to the general texts [28, 5] for broad coverage of these problems.

We do note some rather seminal early literature in the linear setting [6, 37, 16, 39] and the nonlinear setting [6, 27, 1, 3]. The first work on the smoothing topic is the maximum likelihood solution in [6] in both the linear and nonlinear setting. The study of [37] more formally confirms the linear result in [6] and also provides a simpler formulation for the mean and covariance of the smoothing distribution. In the nonlinear setting, the work of [1, 27] introduces an analogue of a type of Kushner-Stratonovich equation (see [5] for this equation in the filtering context) for the smoothing problem. More specifically, [1, 27] propose a deterministic partial differential equation that describes the flow of the smoothing distribution in terms of a backward flow and the standard filtering distribution which acts as the boundary condition (the latter follows from the classical Kushner-Stratonovich equation).

In Section 2 we state the main contribution of this work. Our main result asserts a backward diffusion flow (i.e. a backward stochastic differential equation) whose marginal distribution at any time $0 \leq s \leq t$ is equal to the smoothing distribution $\pi_{t,s}$ when the terminal state is distributed according to the filtering distribution π_t.

This is a novel interpretation of the smoothing solution in terms of a nonlinear diffusion (stochastic) flow (in the spirit of McKean-Vlasov-type processes). This solution contrasts with, and complements, say, the (backward) deterministic flow of probability distributions (viz. a type of Kushner smoothing equation) in [1, 27]. We also provide a number of corollaries of our main result in Section 2.1 including an immediate derivation of the Rauch-Tung-Striebel smoothing equations [37] in the linear setting.

A number of auxiliary contributions are set forth in order to prove our main contribution to the smoothing problem. As is typical, (e.g. see [6, 37, 16, 39, 27, 1, 3, 34]), our smoothing solution requires the formulation of a related filtering problem. In Section 3 we present a brief review of the Kallianpur-Striebel formula. We then provide a novel and more direct approach to deriving weak-versions of the Zakai and the Kushner-Stratonovich equations in Sections 3.1 and 3.2 respectively. We also consider the backward versions of these equations in Section 3.3.

Our approach to the filtering equations in this article combines forward and backward Itô formulas for stochastic transport semigroups with a recent backward version of the Itô-Ventzell formula presented in [13]. This semigroup methodology can be seen as an extension, to the Zakai and Kushner-Stratonovich equations, of the forward-backward stochastic analysis of diffusion flows developed in [10, 11, 13, 23, 24].

Our direct semigroup approach to the forward/backward filtering equations in this work contrasts with classical stochastic partial differential methods and functional analysis in Sobolev spaces; see e.g. the seminal works by Pardoux [31, 32, 34], as well as Krylov and Rozovskii [19, 20]. Related reverse time diffusions and filtered and smoothed densities are also developed in [2, 3] using discrete time approximation techniques, without a detailed discussion on the existence of these densities. We present a number of auxiliary results in this direction throughout Section 3 which are utilised in the proof of our main smoothing result in Section 4.

1.1 Some preliminary notation

This subsection presents some notation needed from the onset.

The signal and the observation defined in (1.1) are column vectors. Unless otherwise stated, we use the letters f and g to denote bounded scalar measurable test functions on some measurable space.

We denote by ∇f the column gradient whenever f is a differentiable function on some Euclidian space, and by $\nabla^2 f$ the Hessian matrix whenever it is twice differentiable.

With $f : \mathbb{R}^m \to \mathbb{R}$, we let $\mathrm{div}_{\alpha_t}(f)$ be the α_t-divergence m-column vector operator with j-th entry given by the formula

$$\mathrm{div}_{\alpha_t}(f)(x)^j := \sum_{1 \leq i \leq m} \partial_{x_i} \left(\alpha_t^{i,j}(x)\ f(x) \right)$$

The generator L_t of the signal X_t is also given by the second order differential operator

$$L_t(f)(x) := \nabla f(x)' b_t(x) + \frac{1}{2}\ \mathrm{Tr}\left(\nabla^2 f(x)\, \alpha_t(x) \right) \quad \text{with the trace operator } \mathrm{Tr}(\cdot).$$

Here and throughout, and without further mention, we assume that functions f acted on by a second-order differential generator are in addition twice differentiable with bounded derivatives.

For a measure μ and test function f of compatible dimension we write,

$$\mu(f) := \int \mu(dx)\, f(x)$$

An integral operator $\mathcal{Q}(x, dz)$ acts on the right on scalar test functions f; and on the left on measures μ according to the formulae,

$$\mathcal{Q}(f)(x) := \int \mathcal{Q}(x, dz)\ f(z) \quad \text{and} \quad (\mu\mathcal{Q})(dz) := \int \mu(dx)\ \mathcal{Q}(x, dz)$$

We extend this operator to an integral operator on matrix functions $h(x) = (h_{i,j}(x))_{i,j}$ by setting,

$$\mathcal{Q}(h)(x)_{i,j} = \mathcal{Q}(h_{i,j})(x)$$

2 Main Result

In further development of this article we assume for any $t > 0$ the conditional distribution π_t has a positive density $p_t := d\pi_t / d\lambda$ w.r.t. the Lebesgue measure λ on $\mathbb{R}^m$. In addition, $p_u(x)$ and its derivative $\nabla p_u(x)$ are uniformly bounded w.r.t. $(u, x) \in [s, t] \times \mathbb{R}^m$, for any given $s > 0$, almost surely w.r.t. the distribution of the observation process. A more detailed discussion on these regularity conditions is provided in Section 2.2.

The main result of the article takes the following form:

Theorem 2.1. *For any $t \geq u \geq s$ we have the transport equation*

$$\pi_{t,s}(dx) = (\pi_{t,u}\, \mathcal{K}_{u,s})(dx) := \int \pi_{t,u}(dz)\ \mathcal{K}_{u,s}(z, dx) \tag{2.1}$$

where $\mathcal{K}_{u,s}$ denotes the Markov semigroup of the backward diffusion flow,

$$d\mathcal{X}_{u,s}(x) = -\left(\left(p_s(\mathcal{X}_{u,s}(x))^{-1}\mathrm{div}_{\alpha_s}(p_s)(\mathcal{X}_{u,s}(x)) - a_s(\mathcal{X}_{u,s}(x))\right)\, ds + \sigma_s(\mathcal{X}_{u,s}(x))\, d\mathcal{W}_s\right) \tag{2.2}$$

with the boundary condition $\mathcal{X}_{u,u}(x) = x$, and where $\mathcal{W}_t \in \mathbb{R}^p$ denotes a p-dimensional Brownian motion independent of the observations.

The proof of the above theorem is provided in Section 4.1. The backward stochastic differential equation (2.2) should be read as shorthand for the backward Itô integration formula,

$$\mathcal{X}_{t,s}(x) = x + \int_s^t \left(p_u(\mathcal{X}_{t,u}(x))^{-1}\mathrm{div}_{\alpha_u}(p_u)(\mathcal{X}_{t,u}(x)) - a_u(\mathcal{X}_{t,u}(x))\right)\, du + \int_s^t \sigma_u(\mathcal{X}_{t,u}(x))\, d\mathcal{W}_u \tag{2.3}$$

with terminal condition $\mathcal{X}_{t,t}(x) = x$. The right-most term in the above formula is an Itô backward stochastic integral such that for any terminal time t this process is a square integrable backward martingale w.r.t. the variable $s \in [0, t]$.

Formally, we may slice the time interval $[s, t]_h := \{u_0, \ldots, u_{n-1}\}$ via some time mesh $u_{i+1} = u_i + h$ from $u_0 = s$ to $u_n = t$ and with time step $h > 0$. In this notation, according to the backward equation (2.2), or (2.3), we compute $\mathcal{X}_{t,u-h}(x)$ from $\mathcal{X}_{t,u}(x)$ using the formula

$$\mathcal{X}_{t,u-h} - \mathcal{X}_{t,u} \ \simeq\ \left(p_u(\mathcal{X}_{t,u})^{-1}\mathrm{div}_{\alpha_u}(p_u)(\mathcal{X}_{t,u}) - a_u(\mathcal{X}_{t,u})\right) h + \sigma_u(\mathcal{X}_{t,u})(\mathcal{W}_u - \mathcal{W}_{u-h}) \tag{2.4}$$

We may provide some comments on the above theorem. By construction, given the observations and for any given $x \in \mathbb{R}^m$ and $t \geq s$, the probability $\mathcal{K}_{t,s}(x, dz)$ introduced in (2.1) coincides with the distribution of the random state $\mathcal{X}_{t,s}(x)$. In addition, for any $t \geq u \geq s$ we have the integral and stochastic semigroup properties,

$$\mathcal{K}_{t,s}(x_2, dx_0) := \int \mathcal{K}_{t,u}(x_2, dx_1)\ \mathcal{K}_{u,s}(x_1, dx_0) \tag{2.5}$$

and

$$\mathcal{X}_{t,s} = \mathcal{X}_{u,s} \circ \mathcal{X}_{t,u} \tag{2.6}$$

where $\mathcal{X}_{u,s} \circ \mathcal{X}_{t,u}$ denotes the composition of the mappings $\mathcal{X}_{u,s}$ and $\mathcal{X}_{t,u}$.

If we let $\mathcal{X}_t$ be a random variable with distribution π_t, for some $t \geq 0$. According to (2.1) the random state $\mathcal{X}_{t,s}(\mathcal{X}_t)$ of the process (2.2) at any given $s \in [0, t]$, is distributed according to $\pi_{t,s} = \pi_t\, \mathcal{K}_{t,s}$. In words, the backward process $\mathcal{X}_{t,s}(\mathcal{X}_t)$ is distributed according to the smoothing distribution $\pi_{t,s}$ for any $s \leq t$ whenever the terminal condition $\mathcal{X}_{t,t}(\mathcal{X}_t) = \mathcal{X}_t$ is distributed according to the filtering distribution π_t. In this sense, (2.2) constitutes a backward nonlinear smoothing diffusion. A forward diffusion flow that has a marginal distribution at any time equal to the filtering distribution is considered in [42, 43].

More generally, we have the backward Itô formula

$$df(\mathcal{X}_{t,s}(x)) = -\mathcal{L}_{s,\pi_s}(f)(\mathcal{X}_{t,s}(x))\ ds - \nabla f(\mathcal{X}_{t,s}(x))'\ \sigma_s(\mathcal{X}_{t,s}(x))\ d\mathcal{W}_s \tag{2.7}$$

with the second order differential operator

$$\mathcal{L}_{s,\pi_s}(f) = \sum_{1 \leq j \leq m} \left(-a_s^j + \frac{1}{p_s}\ \mathrm{div}_{\alpha_s}(p_s)^j\right)\ \partial_{x_j} f + \frac{1}{2} \sum_{1 \leq i,j \leq m} \alpha_s^{i,j}\ \partial_{x_i x_j} f \tag{2.8}$$

Equivalently, we have the backward martingale decomposition

$$f(\mathcal{X}_{t,s}(x)) - f(x) - \int_s^t \mathcal{L}_{u,\pi_u}(f)(\mathcal{X}_{t,u}(x))\, du = \int_s^t \nabla f(\mathcal{X}_{t,u}(x))'\, \sigma_u(\mathcal{X}_{t,u}(x))\, d\mathcal{W}_u \tag{2.9}$$

This yields the backward evolution equations,

$$\partial_s \mathcal{K}_{t,s}(f)(x) = -\mathcal{K}_{t,s}(\mathcal{L}_{s,\pi_s}(f))(x) \tag{2.10}$$

and

$$\partial_s \pi_{t,s}(f) = -\pi_{t,s}(\mathcal{L}_{s,\pi_s}(f)) \tag{2.11}$$

with terminal conditions $\mathcal{K}_{t,t}(f) = f$ and $\pi_{t,t} = \pi_t$. The formula (2.11) coincides with the conditional Fokker-Planck equation in [27], and further developed in [1].

For further discussion on general backward integration of stochastic flows see [10]; see also the appendix of [4] in the context of nonlinear diffusions, and [34] in the context of nonlinear filtering, and [13] on forward-backward perturbation analysis of stochastic flows. Note there is no issue with adaption of the backward process in the sense studied in [35] since we rely only on the independent backward Brownian motion in (2.2). The "backward diffusion" in (2.2) is backward in the sense of a time reversed stochastic differential equation as in [2, 14, 30].

2.1 Some corollaries

We end this introduction with some direct consequences of the above theorem.

Note when $b_t = 0$ the measure π_t coincides with the distribution of the random state X_t of the signal. In this context, $\mathcal{X}_{t,s}(X_t)$ reduces to the time reversal of the signal starting at $\mathcal{X}_{t,t}(X_t) = X_t$ at the terminal time t. Using Theorem 2.1 we recover the fact that the time reversal process of the signal is itself a Markov diffusion [2, 14, 30]. More precisely, we have the corollary:

Corollary 2.2 (Anderson [2])**.** *Assume that $b_t = 0$. For any time horizon $t \geq 0$, the process $\mathfrak{X}_s^t := X_{t-s}$ with $s \in [0,t]$ is a Markov process with generator*

$$\mathfrak{L}_s^t(f) = \sum_{1 \leq j \leq m} \left(\frac{1}{p_{t-s}}\, \mathrm{div}_{\alpha_{t-s}}(p_{t-s})^j - a_{t-s}^j \right) \partial_{x_j} f + \frac{1}{2} \sum_{1 \leq i,j \leq m} \alpha_{t-s}^{i,j}\, \partial_{x_i x_j} f \tag{2.12}$$

We consider now linear-Gaussian filtering/smoothing models with,

$$a_t(x) = A_t\, x, \qquad b_t(x) = B_t\, x \quad \text{and homogeneous diffusion matrix} \quad \sigma_t(x) = \Sigma_t \tag{2.13}$$

for some matrices (A_t, B_t, Σ_t) with appropriate dimensions. Whenever X_0 is a Gaussian random variable with mean $\widehat{X}_0$ and covariance matrix R_0, the optimal filter π_t is a Gaussian distribution with mean $\widehat{X}_t$ and covariance matrix R_t satisfying the Kalman-Bucy and the Riccati equations

$$\left\{ \begin{array}{rcl} d\widehat{X}_t & = & A_t \widehat{X}_t\, dt + R_t B_t' \beta_t^{-1} \left(dY_t - B_t \widehat{X}_t\, dt \right) \\ \partial_t R_t & = & A_t R_t + R_t A_t' + \alpha_t - R_t B_t' \beta_t^{-1} B_t R_t \end{array} \right. \tag{2.14}$$

In this context, we also have that

$$-\, p_s(x)^{-1} \mathrm{div}_{\alpha_s}(p_s)(x) = \alpha_s R_s^{-1} (x - \widehat{X}_s) \tag{2.15}$$

This yields the following corollary:

Corollary 2.3. *For linear Gaussian filtering models (2.13), the diffusion flow $\mathcal{X}_{t,s}(x)$ satisfies the backward evolution equation*

$$d\mathcal{X}_{t,s}(x) = -\left(\left(-A_s\ \mathcal{X}_{t,s}(x) - \alpha_s R_s^{-1}(\mathcal{X}_{t,s}(x) - \widehat{X}_s)\right)\ ds + \Sigma_s\ d\mathcal{W}_s\right) \tag{2.16}$$

with the boundary condition $\mathcal{X}_{t,t}(x) = x$.

Replacing x in (2.16) by a random variable $\mathcal{X}_t$ with distribution π_t for any $t \geq s$ we have that $\mathcal{X}_{t,s}(\mathcal{X}_t)$ has distribution $\pi_{t,s}$. In addition, since the process is linear the distribution $\pi_{t,s}$ is Gaussian with mean $\widehat{X}_{t,s}$ and covariance matrix $R_{t,s}$. The discrete time version of (2.16) can be found in Section 9.9.6 in [12].

Now taking expectations we readily deduce the rather well-known Rauch-Tung-Striebel smoothing equations [37], simplifying the innovation techniques and the sophisticated approximation theory developed in [16, 27, 39], or the formal variational approaches and maximum likelihood techniques presented in the pioneering articles [6, 37].

Corollary 2.4 (Rauch-Tung-Striebel [37])**.** *For any $t \geq s$, the parameters $(\widehat{X}_{t,s}, R_{t,s})$ satisfy the backward evolution equations*

$$\left\{\begin{array}{rcl} \partial_s \widehat{X}_{t,s} & = & A_s \widehat{X}_{t,s} + \alpha_s P_s^{-1}(\widehat{X}_{t,s} - \widehat{X}_s) \\ \partial_s R_{t,s} & = & (A_s + \alpha_s R_s^{-1}) R_{t,s} + R_{t,s}(A_s + \alpha_s R_s^{-1})' - \alpha_s \end{array}\right. \tag{2.17}$$

with terminal conditions $(\widehat{X}_{t,t}, R_{t,t}) = (\widehat{X}_t, R_t)$.

2.2 Comments on our regularity conditions

We end this section with some comments on the regularity conditions discussed at the beginning of Section 2. These conditions are clearly met for linear Gaussian filtering models (see e.g. (2.14) and (2.15)). They are also met for nonlinear models as soon as the signal satisfies a classical controllability-type condition.

Note firstly, whenever the signal is uniformly elliptic, in the sense that $\alpha_t(x) = \sigma_t(x)\sigma_t'(x) \geq \delta\, I$ for some $\delta > 0$, then it is well known that X_t has a smooth positive density w.r.t. the Lebesgue measure on $\mathbb{R}^m$. Nevertheless in many important applications this ellipticity condition is not satisfied. The parabolic Hörmander condition for time varying models [7, 15] is a weaker condition. For linear-Gaussian filtering problems, this condition reduces to the usual controllability condition. Indeed, if we replace the Brownian motions W_t by some arbitrary smooth control functions, all states are accessible from one to another, as soon as the Lie algebra generated by the controlled vector fields is of full rank. This result is also called the Chow-Rashevskii theorem [8, 38]. Under this Hörmander condition, the Hörmander theorem [15] ensures that the signal states have a smooth density w.r.t. the Lebesgue measure on $\mathbb{R}^m$. In addition, for any $s < t$ the Markov transition semigroup $P_{s,t}$ of the signal has a smooth positive density $(x, z) \mapsto p_{s,t}(x, z)$ w.r.t. the Lebesgue measure λ on $\mathbb{R}^m$. In addition, the integral operator $P_{s,t}$ with $s < t$ maps test functions f into bounded smooth functions $P_{s,t}(f)$ given by

$$P_{s,t}(f)(x) = \int P_{s,t}(x, dz)\ f(z) = \int f(z)\ p_{s,t}(x, z)\ dz$$

A natural way to transfer the smoothing properties of $P_{s,t}$ to the optimal filter is to use the following equation

$$\pi_t(f) = \pi_0(P_{0,t}(f)) + \int_0^t \pi_s(P_{s,t}(f)\ (b_s - \pi_s(b_s)))'\ \beta_s^{-1}\ (dY_s - \pi_s(b_s)\ ds) \tag{2.18}$$

given in Theorem 1.1 in [22]. Using this formula we readily check that for any $t > 0$ the conditional distribution π_t has a positive density p_t on $\mathbb{R}^m$. Whenever $\sigma_t(x)$ and $b_t(x)$ are also bounded, Theorem 3.6 in [29] (see also Theorem 6.3 in [23]) also ensures that p_u is smooth, and for any $k \geq 1$, any parameters $h > 0$ and any time horizon $t > 0$ we have

$$\sup_{h \leq s \leq t} \sup_{x \in \mathbb{R}^m} \Big(|p_s(x)| + \|\nabla^k p_s(x)\| \Big) < \infty \tag{2.19}$$

where $\|\cdot\|$ stands for any (equivalent) norm on $\mathbb{R}^m$.

The above estimates are met for linear Gaussian filtering models. Nevertheless, some caution must be used when considering estimates of the form (2.19). Indeed, most of the literature on stochastic partial differential equations arising in nonlinear filtering, such as the strong formulation of the Zakai and Kushner-Stratonovich equations, assume that the sensor function is uniformly bounded, see e.g. [23, 29, 34] and [41, 44]. To the best of our knowledge the extension of the estimate (2.19) to more general unbounded sensor functions is still an open important question.

We also note here that the Kallianpur-Striebel formula [17, 18] is valid as soon as $\beta_u \geq \epsilon\, I$, for some $\epsilon > 0$ and the functions $(a_u(x), b_u(x), \sigma_u(x))$ are smooth with uniformly bounded derivatives w.r.t x of all order on $(u, x) \in [s, t] \times \mathbb{R}^m$, for any $s \leq t$. Weaker conditions can also be found in the book [5] and the recent article [9].

Since X_t has continuous paths, for any continuous function f and any $s \leq t$ the random mapping $u \in [s, t] \mapsto f(X_u)$ is almost surely a uniformly bounded function. In addition, $f(X_t)$ is integrable as soon as f has polynomial growth. Up to some classical localization procedure (see e.g. Chapter 7 in [40]), these rather weak regularity properties also ensure that the integral semigroups that transport (in time) the filtering measures discussed in Section 3, as well as their stochastic partial differential evolution equations, are well defined on any test function with polynomial growth.

3 Nonlinear filtering equations

As is well known (e.g. see [6, 37, 16, 39, 27, 1, 3, 34]), a solution to the smoothing problem will typically make use of the solution of a related filtering in some way. Consequently, we need to present and develop some related filtering results for proving our main result, Theorem 2.1. This section is largely self-contained but it is vital in the proof, in Section 4, of our main result.

The first part of this section presents the classical Kallianpur-Striebel formula which acts as a continuous-time version of Bayes law. In Sections 3.1 and 3.2 respectively we present the Zakai, and Kushner-Stratonovich equations for the flow of the conditional filtering distributions (both unnormalised and normalised). These results are rather well known. For further background on these classical ideas, we refer to the pioneering articles by Kallianpur and Striebel [17, 18], and by Kushner [26] and Zakai [44]. For more recent discussion on these probabilistic models, we refer to [9], and [5, 12], and the references therein. In this article, we present a novel and self contained derivation based on stochastic transport semigroups and their forward evolution equations.

The solution of the Zakai equation is sometimes termed the unnormalized filter. The semigroup that transports these filtering measures (in time) is discussed in Section 3.1; and its normalized version in Section 3.2. Section 3.3 presents a novel direct approach for deriving the backward evolution of these transport semigroups. Our approach in Section 3.3 combines the backward Itô formula for stochastic flows with the backward Itô-Ventzell formula presented in [13].

Now, we introduce some notation/terminology and briefly present the Kallianpur-Striebel formula and the linear semigroup property of unnormalized measures. Let $X_{s,t}(x)$ be the stochastic

flow of the signal on the time interval $[s,t]$ and starting at x at time s. Let $Z_{s,t}(x)$ be the multiplicative functional

$$\log Z_{s,t}(x) := \int_s^t b_u(X_{s,u}(x))' \beta_u^{-1} \; dY_u - \frac{1}{2} \int_s^t b_u(X_{s,u}(x))' \beta_u^{-1} b_u(X_{s,u}(x)) \; du \tag{3.1}$$

When x is replaced by X_s we may write $Z_{s,t}$ instead of $Z_{s,t}(X_s)$, and when $s=0$, we may also write Z_t instead of $Z_{0,t}$. With this notation, we have the classical Kallianpur-Striebel formula,

$$\pi_t(f) = \gamma_t(f)/\gamma_t(1) \quad \text{with} \quad \gamma_t(f) := \mathbb{E}_0\left(f(X_t)\; Z_t\right)$$

Here, $\mathbb{E}_0(\cdot)$ denotes the expectation operator w.r.t. the signal with a fixed observation process.

The transport semigroup of the unnormalized measures γ_t is given for any $s \le t$ by the formula

$$\gamma_t = \gamma_s Q_{s,t} \quad \text{with} \quad Q_{s,t}(f)(x) := \mathbb{E}_0\left(f(X_{s,t}(x))\; Z_{s,t}(x)\right) \tag{3.2}$$

To check this claim observe that,

$$Z_t = Z_s\, Z_{s,t} \quad \Longrightarrow \quad \mathbb{E}_0\left(f(X_t)\; Z_t\right) = \mathbb{E}_0\left(Z_s\; \mathbb{E}_0(f(X_t)\; Z_{s,t} \mid X_s)\right) = \mathbb{E}_0\left(Z_s\; Q_{s,t}(f)(X_s)\right)$$

Now for any $s \le u \le t$ we have

$$\begin{aligned} Q_{s,t}(f)(X_s) \;=\; \mathbb{E}_0(f(X_t)\; Z_{s,t} \mid X_s) \;&=\; \mathbb{E}_0(Z_{s,u}\; \mathbb{E}\left(f(X_t)\;\; Z_{u,t} \mid X_u\right) \mid X_s) \\ &=\; \mathbb{E}_0(Z_{s,u}\; Q_{u,t}(f)(X_u) \mid X_s) \;=\; Q_{s,u}(Q_{u,t}(f))(X_u) \end{aligned}$$

This yields the integral semigroup formula

$$Q_{s,t}(x_0, dx_2) = (Q_{s,u}Q_{u,t})(x_0, dx_2) := \int \; Q_{s,u}(x_0, dx_1)\; Q_{s,u}(x_1, dx_2)$$

In a more compact form, the semigroup property takes the form

$$Q_{s,t} = Q_{s,u}Q_{u,t} \quad \text{with} \quad Q_{t,t} = I \quad \text{where } I \text{ denotes the identity operator.}$$

3.1 Unnormalized stochastic semigroups

Consider the stochastic transport semigroups $\mathbb{P}_{s,t}$ and $\mathbb{Q}_{s,t}$ defined by the composition of functions

$$\mathbb{P}_{s,t}(f)(x) := (f \circ X_{s,t})(x) \quad \text{and} \quad \mathbb{Q}_{s,t}(f)(x) := \mathbb{P}_{s,t}(f)(x)\; Z_{s,t}(x)$$

Using the semigroup properties of the stochastic flow $X_{s,t}(x)$ for any $s \le u \le t$ we check that

$$\mathbb{P}_{s,t}(f)(x) = (f \circ X_{s,t})(x) = (f \circ X_{u,t})(X_{s,u}(x)) = \mathbb{P}_{s,u}(\mathbb{P}_{u,t}(f))(x)$$

Similarly, we have

$$\mathbb{Q}_{s,t}(f)(x) = Z_{s,u}(x)\;\; (Z_{u,t}(X_{s,u}(x))\;\; (f \circ X_{s,t})(X_{s,u}(x))) = \mathbb{Q}_{s,u}(\mathbb{Q}_{u,t}(f))(x)$$

In a more compact form we have the semigroup properties

$$\mathbb{P}_{s,t} = \mathbb{P}_{s,u} \circ \mathbb{P}_{u,t} \quad \text{and} \quad \mathbb{Q}_{s,t} = \mathbb{Q}_{s,u} \circ \mathbb{Q}_{u,t} \quad \text{with} \quad \mathbb{P}_{t,t} = I = \mathbb{Q}_{t,t}$$

Also observe that

$$P_{s,t}(f)(x) := \mathbb{E}_0\left(\mathbb{P}_{s,t}(f)(x)\right) \quad \text{and} \quad Q_{s,t}(f)(x) := \mathbb{E}_0\left(\mathbb{Q}_{s,t}(f)(x)\right)$$

The forward evolution equations of the above semigroups are described in the next proposition.

Proposition 3.1. *For any $t \geq s$ we have the forward stochastic evolution equation*

$$d\mathbb{Q}_{s,t}(f) = \mathbb{Q}_{s,t}(L_t(f))\ dt + \mathbb{Q}_{s,t}(f\, b_t')\ \beta_t^{-1}\ dY_t + \mathbb{Q}_{s,t}((\nabla f)'\ \sigma_t)\ dW_t \tag{3.3}$$

with initial condition $\mathbb{Q}_{s,s}(f) = f$, when $t = s$. In particular, we have the forward equation

$$dQ_{s,t}(f) = Q_{s,t}(L_t(f))\, dt + Q_{s,t}(f\, b_t')\ \beta_t^{-1}\, dY_t \tag{3.4}$$

with the initial condition $Q_{s,s}(f) = f$, when $t = s$.

Proof. Assume that the sensor function $b_u(x)$ is uniformly bounded on $[s,t] \times \mathbb{R}^m$, for any $s \leq t$. In this situation, the random process $(X_{s,u}(x), Z_{s,u}(x))$ also has uniformly bounded absolute moments of any order on any compact interval $[s,t]$, for any time parameters $s \leq t$. In this context, we use Itô formula to check that

$$dZ_{s,t}(x) = Z_{s,t}(x)\ b_t(X_{s,t}(x))'\beta_t^{-1}\ dY_t$$

as well as

$$d\mathbb{P}_{s,t}(f)(x) = \mathbb{P}_{s,t}(L_t(f))(x)\ dt + \mathbb{P}_{s,t}\left(\nabla f'\ \sigma_t\right)(x)\ dW_t$$

An integration by parts yields

$$\begin{aligned} d\mathbb{Q}_{s,t}(f)(x) &= Z_{s,t}(x)\ d\mathbb{P}_{s,t}(f)(x) + \mathbb{P}_{s,t}(f)(x)\ dZ_{s,t}(x) \\ &= \big(L_t(f)(X_{s,t}(x))\ Z_{s,t}(x)\ dt + Z_{s,t}(x)\ f(X_{s,t}(x))\ b_t(X_{s,t}(x))'\beta_t^{-1}\ dY_t\big) \\ &\qquad + Z_{s,t}(x)\ \nabla f(X_{s,t}(x))'\ \sigma_t(X_{s,t}(x))\ dW_t \end{aligned}$$

By classical localization principles of Itô integrals (see for instance Chapter 7 in [40]), the above result is also true for unbounded sensor functions. This ends the proof of (3.3). Taking the expectations, we conclude that

$$d\mathbb{E}_0(\mathbb{Q}_{s,t}(f)(x)) = \mathbb{E}_0\left(\mathbb{Q}_{s,t}(L_t(f))(x)\right)\ dt + \mathbb{E}_0\left(\mathbb{Q}_{s,t}(f\ b_t')(x)\right)\beta_t^{-1}\ dY_t$$

This ends the proof of (3.4). The proof of the proposition is completed. □

Combining (3.2) with Fubini's theorem, we readily check the weak form of Zakai equation given by the formula

$$d\gamma_t(f) = \gamma_t(L_t(f))\, dt + \gamma_t(f\ b_t')\ \beta_t^{-1}\, dY_t \tag{3.5}$$

Arguing as in (2.18), we transfer the smoothing properties of $P_{s,t}$ to $Q_{s,t}$ using the perturbation formulae given for any $s < t$ by

$$Q_{s,t}(f) = P_{s,t}(f) + \int_s^t Q_{s,u}\left(P_{u,t}(f)\ b_u'\right)\ \beta_u^{-1}\, dY_u$$

Arguing as in [44], the above formula shows that for any $s < t$ the integral operator $Q_{s,t}(x_0, dx_1)$ has a density $x_1 \mapsto q_{s,t}(x_0, x_1)$ w.r.t. the Lebesgue measure on $\mathbb{R}^m$ given by the integral equation

$$q_{s,t}(x_0,x_1) = p_{s,t}(x_0,x_1) + \int_s^t \left[\int q_{s,u}(x_0,z)\ p_{u,t}(z,x_1)\ b_u'(z)\ dz\right]\ \beta_u^{-1}\, dY_u \tag{3.6}$$

3.2 Normalized stochastic semigroups

Let $\overline{Z}_{s,t}(x)$ be the multiplicative functional defined as $Z_{s,t}(x)$ by replacing in (3.1) the function b_u and the observation increment dY_u by the centered function $\overline{b}_u$ and the innovation increment $d\overline{Y}_u$ defined by the formulae

$$\overline{b}_u := b_u - \pi_u(b_u) \quad \text{and} \quad d\overline{Y}_u := dY_u - \pi_u(b_u)\ du$$

Under our assumptions, the random process $\pi_t(b_t)$ is almost surely square integrable on any compact time interval so that the innovation process is well defined. Choosing $f = 1$ in (3.5) we check that

$$\log \gamma_t(1) = \int_0^t \pi_u(b_u)' \beta_u^{-1}\ dY_u - \frac{1}{2} \int_0^t \pi_u(b_u)' \beta_u^{-1} \pi_u(b_u)\ du$$

Observe that

$$\pi_s Q_{s,t}(1) = \gamma_t(1)/\gamma_s(1) = \exp\left(\int_s^t \pi_u(b_u)' \beta_u^{-1}\ dY_u - \frac{1}{2} \int_s^t \pi_u(b_u)' \beta_u^{-1} \pi_u(b_u)\ du \right)$$

We also consider the normalized stochastic semigroup

$$\overline{\mathbb{Q}}_{s,t}(f)(x) := (f \circ X_{s,t})(x)\ \overline{Z}_{s,t}(x) = \mathbb{P}_{s,t}(f)(x)\ \overline{Z}_{s,t}(x)$$

Arguing as above, for any $s \leq u \leq t$ we check that

$$\overline{\mathbb{Q}}_{s,t} = \overline{\mathbb{Q}}_{s,u} \circ \overline{\mathbb{Q}}_{u,t} \quad \text{and} \quad \overline{Z}_{s,t}(x) = Z_{s,t}(x)/\pi_s Q_{s,t}(1)$$

Consider the semigroup

$$\overline{Q}_{s,t}(f)(x) := \mathbb{E}_0\left(\overline{\mathbb{Q}}_{s,t}(f)(x)\right) = \mathbb{E}_0\left(f(X_{s,t}(x))\ \overline{Z}_{s,t}(x)\right) = Q_{s,t}(f)(x)/\pi_s Q_{s,t}(1)$$

In this notation, using the same arguments as in the proof of Proposition 3.1 we have the following forward evolution equations.

Proposition 3.2. *For any given time horizon s and for any $t \geq s$ we have the forward stochastic evolution equation*

$$d\overline{\mathbb{Q}}_{s,t}(f) = \overline{\mathbb{Q}}_{s,t}(L_t(f))\ dt + \overline{\mathbb{Q}}_{s,t}(f\, \overline{b}_t')\ \beta_t^{-1}\ d\overline{Y}_t + \overline{\mathbb{Q}}_{s,t}((\nabla f)'\ \sigma_t)\ dW_t$$

with initial condition $\overline{Q}_{s,s}(f) = f$, when $t = s$. In particular, we have the forward equation

$$d\overline{Q}_{s,t}(f) = \overline{Q}_{s,t}(L_t(f))\, dt + \overline{Q}_{s,t}(f\ \overline{b}_t')\ \beta_t^{-1}\, d\overline{Y}_t$$

with the initial condition $\overline{Q}_{s,s}(f) = f$, when $t = s$.

The above proposition yields the weak form of the Kushner-Stratonovich equation defined by

$$d\pi_t(f) = \pi_t(L_t(f))\, dt + \pi_t(f\ \overline{b}_t)'\ \beta_t^{-1}\, d\overline{Y}_t \tag{3.7}$$

Formally, using the same notation as in (3.11) we have the forward approximation equation

$$\pi_{u+h}(f) \simeq \pi_u(f) + \pi_u(L_u(f))\, h + \pi_u(f\ \overline{b}_u)'\ \beta_u^{-1}\ \left(\overline{Y}_{u+h} - \overline{Y}_u\right) \tag{3.8}$$

3.3 Backward evolution equations

This section is concerned with the backward evolution equation associated with the unnormalized semigroup $\mathbb{Q}_{s,t}$ and its normalized version. The main result of this section is the following theorem:

Theorem 3.3. *For any twice differentiable function f with bounded derivatives and for any $s \leq t$ we have the backward evolution equation*

$$\begin{aligned} d\mathbb{Q}_{s,t}(f)(x) &= -\left(\nabla \mathbb{Q}_{s,t}(f)(x)' \; a_s(x) + \frac{1}{2} \operatorname{Tr}\left(\nabla^2 \mathbb{Q}_{s,t}(f)(x) \; \alpha_s(x)\right)\right) ds \\ &\quad - \mathbb{Q}_{s,t}(f)(x) \; b_s(x)' \beta_s^{-1} \; dY_s - \nabla \mathbb{Q}_{s,t}(f)(x)' \; \sigma_s(x) \; dW_s \end{aligned} \tag{3.9}$$

with terminal condition $\mathbb{Q}_{t,t}(f) = f$, when $s = t$. In particular, we have the backward equation

$$dQ_{s,t}(f) = -\left(L_s(Q_{s,t}(f))\, ds + Q_{s,t}(f) \; b_s' \; \beta_s^{-1} \, dY_s\right) \tag{3.10}$$

with terminal condition $Q_{t,t}(f) = f$, when $s = t$.

Proof. We use a direct approach combining the backward filtering calculus developed in [23, 41] based on the backward Itô calculus developed in [10, 11, 21, 25], see also the more recent article [13] and references therein.

Consider the stochastic flow $\chi_{s,t}(\overline{x})$ starting at

$$\chi_{s,s}(\overline{x}) = \overline{x} := \begin{pmatrix} x \\ z \end{pmatrix} \in (\mathbb{R}^m \times \mathbb{R})$$

on the time interval $[s, \infty[$ and given for any $t \geq s$ by

$$\chi_{s,t}(\overline{x}) := \begin{pmatrix} X_{s,t}(x) \\ Z_{s,t}(x) \; z \end{pmatrix} \in (\mathbb{R}^m \times \mathbb{R})$$

We set

$$\mathcal{B}_t(\overline{x}) := \begin{pmatrix} a_t(x) \\ 0 \end{pmatrix} \qquad \mathcal{U}_t := \begin{pmatrix} W_t \\ Y_t \end{pmatrix}$$

$$\Lambda_t(\overline{x}) := \begin{pmatrix} \sigma_t(x) & 0 \\ 0 & z \; b_t(x)' \beta_t^{-1} \end{pmatrix} \quad \text{and} \quad \mathcal{A}_t(\overline{x}) := \Lambda_t(\overline{x}) \, \Lambda_t(\overline{x})'$$

Assume that the sensor function $b_u(x)$ is uniformly bounded on $[s,t] \times \mathbb{R}^m$, for any $s \leq t$. Then, the process $(Z_{s,u}(x), \chi_{s,u}(\overline{x}))$ has continuous partial derivatives and also has uniformly bounded absolute moments of any order on $([s,t] \times \mathbb{R}^m)$, for any $s \leq t$. In this situation, we have the forward stochastic evolution equation

$$d\chi_{s,t}(\overline{x}) = \mathcal{B}_t(\chi_{s,t}(\overline{x})) \; dt + \Lambda_t(\chi_{s,t}(\overline{x})) \; d\mathcal{U}_t$$

For any twice differentiable function F on $(\mathbb{R}^m \times \mathbb{R})$ with bounded derivatives we also have the backward equation

$$\begin{aligned} &d(F \circ \chi_{s,t})(\overline{x}) \\ &= -\left(\nabla (F \circ \chi_{s,t})(\overline{x})' \; \mathcal{B}_s(\overline{x}) + \frac{1}{2} \operatorname{Tr}\left(\nabla^2 (F \circ \chi_{s,t})(\overline{x}) \; \mathcal{A}_s(\overline{x})\right)\right) ds - \nabla (F \circ \chi_{s,t})(\overline{x})' \; \Lambda_s(\overline{x}) \, d\mathcal{U}_s \end{aligned}$$

A proof of the above formula can be found in the articles [10, 11], see also [13]. Choosing the function $F(\overline{x}) = f(x)\ z$, for some twice differentiable function f on $\mathbb{R}^m$ with bounded derivatives and letting $z = 1$ we check that

$$
\begin{aligned}
& d(f(X_{s,t}(x))Z_{s,t}(x)) \\
& = -\left(\nabla(f(X_{s,t}(x))Z_{s,t}(x))'\ a_s(x) + \frac{1}{2}\ \mathrm{Tr}\left(\nabla^2(f(X_{s,t}(x))Z_{s,t}(x))\ \alpha_s(x)\right)\right)\ ds \\
& \qquad - (f(X_{s,t}(x))Z_{s,t}(x))\ b_s(x)'\beta_s^{-1}\ dY_s - \nabla(f(X_{s,t}(x))Z_{s,t}(x))'\ \sigma_s(x)\ dW_s
\end{aligned}
$$

This ends the proof of (3.9). By localization arguments, the above result is also true for unbounded sensor functions. Integrating the flow of the signal we obtain (3.10). This ends the proof of the theorem. □

We can also check (3.10) considering a discrete time interval $[s,t]_h := \{t_0, \ldots, t_{n-1}\}$ associated with some refining time mesh $t_{i+1} = t_i + h$ from $t_0 = s$ to $t_n = t$, for some time step $h > 0$. By (3.4), for any $u \in [s,t]_h$ we compute $Q_{u,t}(f)$ from $Q_{u+h,t}(f)$ using the backward equation

$$
\begin{aligned}
Q_{u,t}(f) &= Q_{u+h,t}(f) + (Q_{u,u+h} - I)\,(Q_{u+h,t}(f)) \\
&\simeq Q_{u+h,t}(f) + L_u(Q_{u+h,t}(f))\,h + Q_{u+h,t}(f)\ b'_u\ \beta_u^{-1}\,(Y_{u+h} - Y_u)
\end{aligned}
\tag{3.11}
$$

For null sensor functions the evolution equation (3.9) coincides with the backward Itô formula discussed in [10, 11, 13, 23, 24].

Choosing $f = 1$ in (3.10) we recover the backward evolution of the likelihood function presented in [3, 33] (see formula (5.9) in [3] and equation (3.15) in [33]). Arguing as in (3.6), using (3.10) we check the perturbation formulae given for any $s < t$ by,

$$
Q_{s,t}(f) = P_{s,t}(f) + \int_s^t P_{s,u}\left(Q_{u,t}(f)\ b'_u\right)\ \beta_u^{-1}\, dY_u
$$

Thus, for any $s < t$ the integral operator $Q_{s,t}(x_0, dx_1)$ has a density $(x_0, x_1) \mapsto q_{s,t}(x_0, x_1)$ given by (3.6) and the integral formula,

$$
q_{s,t}(x_0, x_1) = p_{s,t}(x_0, x_1) + \int_s^t \left[\int p_{s,u}(x_0, z)\ q_{u,t}(z, x_1)\ b'_u(z)\ dz\right]\ \beta_u^{-1}\, dY_u \tag{3.12}
$$

Using the same arguments as in the proof of Theorem 3.3 we also have the following backward evolution equation.

Proposition 3.4. *For any twice differentiable function f with bounded derivatives and for any $s \leq t$ we also have the backward equation*

$$
\begin{aligned}
d\overline{\mathbb{Q}}_{s,t}(f)(x) &= -\left(\nabla\,\overline{\mathbb{Q}}_{s,t}(f)(x)'\ a_s(x) + \frac{1}{2}\ \mathrm{Tr}\left(\nabla^2\,\overline{\mathbb{Q}}_{s,t}(f)(x)\ \alpha_s(x)\right)\right)\ ds \\
&\qquad - \overline{\mathbb{Q}}_{s,t}(f)(x)\ \overline{b}_s(x)'\beta_s^{-1}\ d\overline{Y}_s - \nabla\overline{\mathbb{Q}}_{s,t}(f)(x)'\ \sigma_s(x)\ dW_s
\end{aligned}
$$

with terminal condition $\overline{\mathbb{Q}}_{t,t}(f) = f$. In particular, we have the backward equation,

$$
d\overline{Q}_{s,t}(f) = -\left(L_s(\overline{Q}_{s,t}(f))\,ds + \overline{Q}_{s,t}(f)\ \overline{b}'_s\ \beta_s^{-1}\, d\overline{Y}_s\right) \tag{3.13}
$$

with terminal condition $\overline{Q}_{t,t}(f) = f$.

Using the same notation as in (3.11), we also have the approximating backward equation

$$
\overline{Q}_{u,t}(f) \simeq \overline{Q}_{u+h,t}(f) + L_u(\overline{Q}_{u+h,t}(f))\,h + \overline{Q}_{u+h,t}(f)\ \overline{b}'_u\ \beta_u^{-1}\ \left(\overline{Y}_{u+h} - \overline{Y}_u\right) \tag{3.14}
$$

4 Smoothing semigroups and proof of the main result

This section is concerned with forward-backward evolution equations for the conditional smoothing distribution and the proof of our main result.

Let $\mathcal{K}_{t,s}$ be the backward integral operator defined by,

$$\mathcal{K}_{t,s}(f)(x) := \int \pi_s(dz)\ \frac{d\overline{Q}_{s,t}(z,\cdot)}{d\pi_t}(x)\, f(z) \tag{4.1}$$

For any $s \leq u \leq t$ we have the backward semigroup property,

$$\mathcal{K}_{t,s} = \mathcal{K}_{t,u}\, \mathcal{K}_{u,s} \tag{4.2}$$

which follows via,

$$\begin{aligned}(\mathcal{K}_{t,u}\, \mathcal{K}_{u,s})(f)(x) &= \int \pi_s(dx_0)\ \overline{Q}_{s,u}(x_0, dx_1)\ \frac{d\overline{Q}_{u,t}(x_1,\cdot)}{d\pi_t}(x)\, f(x_0) \\ &= \int \pi_s(dx_0)\ \frac{d\overline{Q}_{s,t}(x_0,\cdot)}{d\pi_t}(x)\, f(x_0) \ = \ \mathcal{K}_{t,s}(f)(x)\end{aligned}$$

and where we exploit the semigroup properties of the operators $\overline{Q}_{s,t}$.

Also observe that for any $t > s > 0$ the integral operator $\mathcal{K}_{t,s}(x_1, dx_0)$ has a density $(x_1, x_0) \mapsto k_{s,t}(x_1, x_0)$ w.r.t. the Lebesgue measure on $\mathbb{R}^m$ given by,

$$k_{t,s}(x_1, x_0) := p_s(x_0)\ \overline{q}_{s,t}(x_0, x_1)/p_t(x_1) \quad \text{with} \quad \overline{q}_{s,t}(x_0, x_1) = q_{s,t}(x_0, x_1)/\pi_s(Q_{s,t}(1))$$

The function $q_{s,t}$ denotes the density of the integral operator $Q_{s,t}$ discussed in (3.6) and (3.12).

Now, for any pair of functions (f, g) we readily check the duality formula,

$$\pi_s\left(f\, \overline{Q}_{s,t}(g)\right) = \pi_t\left(\mathcal{K}_{t,s}(f)\, g\right) \tag{4.3}$$

The following technical result is key in the proof of Theorem 2.1.

Lemma 4.1. *For any time parameter $s \leq t$ we have the forward-backward differential equation*

$$\partial_s\left(\pi_s\left(f\ \overline{Q}_{s,t}(g)\right)\right) = -\pi_s\left(\overline{Q}_{s,t}(g)\ \mathcal{L}_{s,\pi_s}(f)\right) \tag{4.4}$$

with the second order differential operator

$$\mathcal{L}_{s,\pi_s}(f) \ := \ -L_s(f) + \frac{1}{p_s} \sum_{1\leq i,j\leq m} \partial_{x_i}\left(p_s\ \alpha_s^{i,j}\ \partial_{x_j} f\right)$$

Proof. Observe that (4.4) does not involve the derivatives of the function g. Thus, up to a smooth mollifier's type approximation of the function g, it suffices to check (4.4) for any pair of bounded and twice differentiable functions f, g with bounded differentials. Arguing as in the proof of Proposition 3.1 and Theorem 3.3, it suffices to prove the result for uniformly bounded sensor functions $b_u(x)$ on $[s,t] \times \mathbb{R}^m$, for any $s \leq t$.

In this situation, for any time horizon t, combining the Kushner-Stratonovich equation (3.7) with the backward equation (3.13) for any $s \leq t$, we check the forward-backward evolution equation

$$\partial_s\left(\pi_s\left(f\ \overline{Q}_{s,t}(g)\right)\right) = \pi_s\left(L_s(f\ \overline{Q}_{s,t}(g)) - f\ L_s(\overline{Q}_{s,t}(g))\right) \tag{4.5}$$

The above equation can be proved using the backward Itô-Ventzell formula in [13]. We use the same notation as in the proof of Theorem 3.3. Let $\overline{Z}_{s,t}(x)$ be the multiplicative functional defined as $Z_{s,t}(x)$ by replacing the function b_u and the observation Itô-increment dY_u by the centered function $\overline{b}_u$ and the innovation increment $d\overline{Y}_u$.

Consider the backward random field $F_{s,t}$ with terminal condition $F_{t,t}(\overline{x}) = f(x)g(x)z$ defined by the formula

$$F_{s,t}(\overline{x}) := f(x)\ \overline{\mathbb{Q}}_{s,t}(g)(x)\ z \quad \text{and we set} \quad \overline{\chi}_s := \begin{pmatrix} X_s \\ \overline{Z}_s \end{pmatrix} \in (\mathbb{R}^m \times \mathbb{R}).$$

In this notation, we have

$$\mathbb{E}_0\left(F_{s,t}(\overline{\chi}_s)\right) = \mathbb{E}_0\left(f(X_s)\ \overline{Z}_s\ \mathbb{E}_0\left(\overline{\mathbb{Q}}_{s,t}(g)(X_s)|\ (X_s, Z_s)\right)\right) = \pi_s(f\ \overline{Q}_{s,t}(g))$$

Observe that

$$F_{s,t}(\overline{x}) = f(x)\ (F \circ \overline{\chi}_{s,t})(\overline{x})$$

with the function

$$F(\overline{x}) := g(x)\ z \quad \text{and the stochastic flow} \quad \overline{\chi}_{s,t}(x,z) := \begin{pmatrix} X_{s,t}(x) \\ \overline{Z}_{s,t}(x)\ z \end{pmatrix}$$

Following the proof of Theorem 3.3, we check that

$$dF_{s,t}(\overline{x}) = f(x)\ d(F \circ \overline{\chi}_{s,t})(\overline{x}) = -\left(\mathcal{G}_{s,t}(\overline{x})\ ds + \mathcal{H}_{s,t}(\overline{x})\ d\mathcal{U}_s\right)$$

with the drift function

$$\mathcal{G}_{s,t}(\overline{x}) := \ f(x)\ z\ \left(\nabla\,\overline{\mathbb{Q}}_{s,t}(g)(x)'\ a_s(x) + \frac{1}{2}\ \mathrm{Tr}\left(\nabla^2\overline{\mathbb{Q}}_{s,t}(g)(x)'\ \alpha_s(x)\right)\right)$$

and the diffusion term

$$\mathcal{H}_{s,t}(\overline{x})\ d\mathcal{U}_s := \ f(x)\ z\ \left(\nabla\,\overline{\mathbb{Q}}_{s,t}(g)(x)'\ \sigma_s(x)\ dW_s + \overline{\mathbb{Q}}_{s,t}(g)(x)\ b_s(x)'\beta_s^{-1}\ dY_s\right)$$

Applying the backward Itô-Ventzell formula [13] we check that

$$dF_{s,t}(\overline{\chi}_s) = (dF_{s,t})(\overline{\chi}_s) + \nabla F_{s,t}(\overline{\chi}_s)'\ d\chi_s + \frac{1}{2}\ \mathrm{Tr}\left(\nabla^2 F_{s,t}(\chi_s)'\ \mathcal{A}_t(\overline{\chi}_s)\right) ds$$

from which we conclude that

$$\begin{aligned} dF_{s,t}(\overline{\chi}_s) &= \overline{Z}_s\left(\nabla\left(\overline{\mathbb{Q}}_{s,t}(g)(x)\ f(x)\right)'_{|x=X_s} - f(X_s)\ \overline{Z}_s\ \nabla\,\overline{\mathbb{Q}}_{s,t}(g)(X_s)'\right)\sigma_s(X_s)\,dW_s \\ &- f(X_s)\overline{Z}_s\ \left(\nabla\,\overline{\mathbb{Q}}_{s,t}(g)(X_s)'\ a_s(X_s)\ ds + \frac{1}{2}\ \mathrm{Tr}\left(\nabla^2\,\overline{\mathbb{Q}}_{s,t}(g)(X_s)\ \alpha_s(X_s)\right)\right) ds \\ &+ \overline{Z}_s\left(\nabla\left(\overline{\mathbb{Q}}_{s,t}(g)(x)\ f(x)\right)'_{|x=X_s}\ a_s(X_s)\ ds + \ \frac{1}{2}\ \mathrm{Tr}\left(\nabla^2\left(\overline{\mathbb{Q}}_{s,t}(g)(x)\ f(x)\right)'_{|x=X_s}\ \alpha_s(X_s)\right)\right) ds \end{aligned}$$

We end the proof of (4.5) by simple integration.

To take the final step, we recall the integration by parts formula

$$L_t(fg) = f\ L_t(g) + g\ L_t(f) + \Gamma_{L_t}(f,g)$$

with the carré-du-champ (a.k.a. square field) operator Γ_{L_t} associated with the generator L_t defined by

$$\Gamma_{L_t}(f,g) := (\nabla f)' \alpha_t \nabla g$$

Combining (4.5) with the above formula we check that

$$\partial_s \left(\pi_s \left(f \ \overline{Q}_{s,t}(g) \right) \right) = \pi_s \left(L_s(f) \ \overline{Q}_{s,t}(g) \right) + \pi_s \left(\Gamma_{L_s} \left(\overline{Q}_{s,t}(g), f \right) \right)$$

On the other hand, by an integration by parts we have

$$\pi_s \left(\Gamma_{L_s} \left(\overline{Q}_{s,t}(g), f \right) \right) = - \sum_{i,j} \int \ p_s(x) \ \overline{Q}_{s,t}(g)(x) \frac{1}{p_s(x)} \partial_{x_i} \ \left(p_s(x) \ \alpha_t^{i,j} \ \partial_{x_j} f(x) \right) dx$$

This ends the proof of the lemma. □

Another approach for finding (4.5) is to use for any $u \in [s,t]_h$ the decomposition

$$\begin{aligned} & \pi_{u+h} \left(f \, \overline{Q}_{u+h,t}(g) \right) - \pi_u \left(f \, \overline{Q}_{u,t}(g) \right) \\ & \quad = \pi_u \left(f \left(\overline{Q}_{u+h,t} - \overline{Q}_{u,t} \right)(g) \right) + \left(\pi_{u+h} - \pi_u \right) \left(f \ \overline{Q}_{u+h,t}(g) \right) \end{aligned} \tag{4.6}$$

Note that π_u depends on the observations $(Y_s - Y_0)$ from $s = 0$ up to time $s = u$, while the increment $\overline{Q}_{u,t}$ is computed backward in time and only depends on the observations $(Y_s - Y_u)$ from $s > u$ up to $s = t$. Conversely, π_{u+h} depends on the observations $(Y_s - Y_0)$ from $s = 0$ up to time $s = u + h$, while $\overline{Q}_{u+h,t}$ is computed backward in time and only depends on the observations $(Y_s - Y_{u+h})$ from $s > u + h$, up to time $s = t$.

Following the two-sided stochastic integration calculus developped by Pardoux and Protter in [36] (see also [13] for extended versions to interpolating stochastic flows), combining the forward (3.8) with the backward equation (3.14), when $h \simeq 0$ we can check the approximation,

$$\sum_{u \in [s,t]_h} \left\{ \pi_{u+h} \left(f \ \overline{Q}_{u+h,t}(g) \right) - \pi_u \left(f \ \overline{Q}_{u,t}(g) \right) - \pi_u \left(L_u(f \ \overline{Q}_{u+h,t}(g)) - f \ L_u(\overline{Q}_{u+h,t}(g)) \right) \ h \right\} \simeq \ 0$$

4.1 Proof of Theorem 2.1

With the definition in (4.1) we have,

$$\pi_{t,s}(dx) = (\pi_t \, \mathcal{K}_{t,s})(dx) = \pi_s(dx) \, \overline{Q}_{s,t}(1)(x) \tag{4.7}$$

The formulation of the conditional distribution $\pi_{t,s}$ of X_s given $\mathcal{Y}_t$ in (4.7) is rather well known, see e.g. Theorem 3.7 and Corollary 3.8 in [34], as well as equation (3.9) in [3]. The proof of this formula is a direct consequence of (4.1). With (4.2) we have,

$$\pi_t \mathcal{K}_{t,s} = \pi_{t,u} \mathcal{K}_{u,s} = \pi_{t,s}$$

Thus with $\mathcal{K}_{t,s}$ as defined in (4.1) we immediately have the transport equation in (2.1).

It remains to show that this integral operator (as defined in (4.1)) is also the Markov transition kernel of the backward diffusion flow in (2.2). The rest of the proof of Theorem 2.1 is a consequence of the duality formula (4.3) and Lemma 4.1.

Rewritten in a slightly different form, the duality formula (4.3) reads as follows,

$$\mathbb{E} \left(f(X_s) \ g(X_t) \mid \mathcal{Y}_t \right) = \mathbb{E} \left(\mathcal{K}_{t,s}(f)(X_t) \ g(X_t) \mid \mathcal{Y}_t \right)$$

This implies that

$$\mathcal{K}_{t,s}(f)(X_t) = \mathbb{E}\left(f(X_s) \mid X_t,\ \mathcal{Y}_t\right)$$

Finally, combining (4.4) with the duality formula (4.3) we have

$$\pi_t\left(g\ \partial_s \mathcal{K}_{t,s}(f)\right) = -\pi_t\left(g\ \mathcal{K}_{t,s}(\mathcal{L}_{s,\pi_s}(f))\right)$$

Since the above formula is valid for any test function g and π_t has a bounded positive density, we check the backward Kolmogorov equation

$$\partial_s \mathcal{K}_{t,s}(f)(x) = -\mathcal{K}_{t,s}(\mathcal{L}_{s,\pi_s}(f))(x) \tag{4.8}$$

with terminal condition $\mathcal{K}_{t,t}(f) = f$, when $s = t$, for almost every $x \in \mathbb{R}^m$ (and almost surely w.r.t. the law of the observation process from the origin up to the time t). Since both terms in (4.8) are continuous, the above equality holds for any $x \in \mathbb{R}^m$, almost surely.

We now complete the proof by showing that the integral operator $\mathcal{K}_{t,s}(x, dz)$ (defined in (4.1)) does indeed coincide with the transition kernel associated with the flow $\mathcal{X}_{t,s}(x)$ in (2.2). Firstly, observe that (4.8) coincides with the backward Kolmogorov equation (2.11) associated with the transition semigroup of the stochastic flow $\mathcal{X}_{t,s}(x)$. Denote this transition semigroup by $\overline{\mathcal{K}}_{t,s}(x, dz)$ temporarily.

By the semigroup properties of $\overline{\mathcal{K}}_{t,s}$, for any $s \leq u \leq t$ and any smooth function f we have

$$\partial_u \overline{\mathcal{K}}_{t,s}(f) = 0 = \partial_u(\overline{\mathcal{K}}_{t,u}(\overline{\mathcal{K}}_{u,s}(f))) = -\overline{\mathcal{K}}_{t,u}(\mathcal{L}_{u,\pi_u}(\overline{\mathcal{K}}_{u,s}f)) + \overline{\mathcal{K}}_{t,u}(\partial_u \overline{\mathcal{K}}_{u,s}(f))$$

Choosing $u = t$ we obtain the forward equation

$$\partial_t \overline{\mathcal{K}}_{t,s}(f) = \mathcal{L}_{t,\pi_t}(\overline{\mathcal{K}}_{t,s}(f))$$

Arguing as above, this implies that

$$\partial_u(\mathcal{K}_{t,u}(\overline{\mathcal{K}}_{u,s}(f))) = -\mathcal{K}_{t,u}(\mathcal{L}_{u,\pi_u}(\overline{\mathcal{K}}_{u,s}f)) + \mathcal{K}_{t,u}(\mathcal{L}_{u,\pi_u}(\overline{\mathcal{K}}_{u,s}(f))) = 0$$

Integrating over the interval $[s, t]$ we check that $\mathcal{K}_{t,s} = \overline{\mathcal{K}}_{t,s}$. This ends the proof of Theorem 2.1. □

References

[1] B.D.O. Anderson. Fixed interval smoothing for nonlinear continuous time systems. Information and Control. vol. 20. pp. 294-300 (1972).

[2] B.D.O. Anderson. Reverse-time diffusion equation models. Stochastic Processes and their Applications. vol.12, no. 3. pp. 313–326 (1982).

[3] B.D.O. Anderson and I.B. Rhodes. Smoothing Algorithms for Nonlinear Finite-Dimensional Systems. Stochastics. vol. 9, no. 1-2. pp. 139–165 (1983).

[4] M. Arnaudon and P. Del Moral. A variational approach to nonlinear and interacting diffusions. Stochastic Analysis and Applications. vol. 37, no. 5. pp. 717–748 (2019).

[5] A. Bain and D. Crisan. Fundamentals of Stochastic Filtering. Springer (2009).

[6] A.E. Bryson and M. Frazier. Smoothing for linear and nonlinear dynamic systems. In Proceedings of the Optimum System Synthesis Conference. Technical Document Report No.: ASD-TDR-63-119, Aeronautical Systems Division, Wright-Patterson Air Force Base, Ohio. pp. 353–364 (1963).

[7] P. Cattiaux and L. Mesnager. Hypoelliptic non-homogeneous diffusions. Probability Theory and Related Fields. vol. 123, no. 4, pp. 453–483 (2002).

[8] W.L. Chow. Über Systeme von linearen partiellen Differentialgleichungen erster Ordnung, Mathematische Annalen. vol. 117. pp. 98–105 (1939).

[9] T. Cass, M. Clark, and D. Crisan. The Filtering Equations Revisited. In Proc. of the Stochastic Analysis and Applications. pp. 129–162 (2014).

[10] G. Da Prato, J.L. Menaldi, and L. Tubaro. Some results of backward Ito formula. Stochastic Analysis and Applications. vol. 25, no. 3. pp. 679–703 (2007).

[11] G. Da Prato. Some remarks about backward Itô formula and applications. Stochastic Analysis and Applications. vol. 16, no. 6. pp. 993–1003 (1998).

[12] P. Del Moral and S. Penev. Stochastic Processes: From Applications to Theory. CRC Press (2017).

[13] P. Del Moral and S.S. Singh. A forward-backward stochastic analysis of diffusion flows. arXiv e-print, `arXiv:1906.09145` (2019).

[14] H.G. Haussmann and E. Pardoux. Time reversal of diffusions. The Annals of Probability. vol. 14, no. 4. pp. 1188–205 (1986).

[15] L. Hörmander. Hypoelliptic second order differential equations. Acta Mathematica. vol. 119. pp. 147–171 (1967).

[16] T. Kailath, P. Frost. An innovations approach to least-squares estimation - Part II: Linear smoothing in additive white noise. IEEE Transactions on Automatic Control. vol. 13, no. 6. pp. 655–660 (1968).

[17] G. Kallianpur and C. Striebel. Estimation of stochastic systems: arbitrary system process with additive noise observation errors. The Annals of Mathematical Statistics. vol. 39, no. 3. pp. 785–801 (1968).

[18] G. Kallianpur and C. Striebel. Stochastic differential equations occurring in the estimation of continuous parameter stochastic processes. Theory of Probability & Its Applications. vol. 14, no. 4. pp. 567–594 (1969).

[19] N.V. Krylov and B.L. Rozowskii. On the Cauchy problem for linear stochastic partial differential equations. Mathematics of the USSR-Izvestiya. vol. 11, no. 4. (1977).

[20] N.V. Krylov and B.L. Rozowskii. Conditional distributions of diffusion processes. Mathematics of the USSR-Izvestiya. vol. 12, no. 2. (1978).

[21] N.V. Krylov and B.L. Rozowskii. On the first integrals and Liouville equations for diffusion processes. In Stochastic Differential Systems: Proc. 3rd SFSP-WG 7/1. Part of the Lecture Notes in Control and Information Sciences. vol. 36. pp. 117-125. Springer (1981).

[22] H. Kunita. Asymptotic behavior of the nonlinear filtering errors of Markov processes. Journal of Multivariate Analysis. vol. 1, no. 4. pp 365–393 (1971).

[23] H. Kunita. Stochastic partial differential equations connected with non-linear filtering. In Nonlinear Filtering and Stochastic Control. Part of the Lecture Notes in Mathematics. vol. 972. pp. 100–169. Springer (1982).

[24] H. Kunita. On backward stochastic differential equations. Stochastics. vol. 6, no. 3-4. pp. 293–313 (1982).

[25] H. Kunita. First order stochastic partial differential equations. In Stochastic Analysis: Proc. of the Taniguchi International Symposium on Stochastic Analysis. pp. 249–269 (1984).

[26] H.J. Kushner. On the differential equations satisfied by conditional probability densities of Markov processes, with applications. Journal of the Society for Industrial and Applied Mathematics, Series A: Control. vol. 2, no. 1. pp. 106–119 (1964).

[27] C.T. Leondes, J.B. Peller, and E.B. Stear. Nonlinear smoothing theory. IEEE Transactions on Systems Science and Cybernetics. vol. 6, no. 1. pp. 63–71 (1970).

[28] J.S. Meditch. A Survey of Data Smoothing for Linear and Nonlinear Dynamic Systems. Automatica. vol. 9. pp. 151–162 (1973).

[29] D. Michel. Régularité des lois conditionnelles en théorie du filtrage non-linéaire et calcul des variations stochastique. Journal of Functional Analysis. vol. 41, no. 1. pp. 8–36 (1981).

[30] A. Millet, D. Nualart, and M. Sanz. Integration by parts and time reversal for diffusion processes. The Annals of Probability. vol. 17, no.1. pp. 208–238 (1989).

[31] E. Pardoux. Equations aux dérivées partielles stochastiques non linéaires monotones. Doctoral Thesis, Université Paris XI, Orsay (1975).

[32] E. Pardoux. Stochastic partial differential equations and filtering of diffusion processes. Stochastics. vol. 3, no. 1-4. pp. 127–167 (1980).

[33] E. Pardoux, Non-linear filtering, prediction and smoothing, in M. Hazewinkel and J. C. Willems (eds.), Stochastic Systems: The Mathematics of Filtering and Identijication and Applications, Reidel, Dordrecht (1981).

[34] E. Pardoux. Équations du filtrage non linéaire de la prédiction et du lissage. Stochastics. vol. 6, no. 3–4. pp. 193–231 (1982).

[35] E. Pardoux and S. Peng. Adapted solution of a backward stochastic differential equation. Systems & Control Letters. vol. 14, no. 1. pp. 55–61 (1990).

[36] E. Pardoux and P. Protter. A two-sided stochastic integral and its calculus. Probability Theory and Related Fields. vol. 76, no. 1. pp. 15–49 (1987).

[37] H.E. Rauch, F. Tung, and C.T. Striebel. Maximum likelihood estimates of linear dynamic systems. AIAA Journal. vol. 3, no. 8. pp. 1445–1450 (1965).

[38] P.K. Rashevskii. About connecting two points of complete non-holonomic space by admissible curve (in Russian). Uch. Zapiski Ped. Instit. Libknexta. vol. 2. pp 83–94 (1938).

[39] M. Rutkowski. A simple proof for the Kalman-Bucy smoothed estimate formula. Statistics & Probability Letters. vol. 17, no. 5. pp. 377–385 (1993).

[40] J. M. Steele. Stochastic Calculus and Financial Applications. Springer (2012).

[41] A.Y. Veretennikov. On backward filtering equations for SDE systems (direct approach). In Stochastic Partial Differential Equations: Part of the Lecture Notes Series. vol. 216. pp. 304–311. London Mathematical Society (1995).

[42] T. Yang, P.G. Mehta, and S.P. Meyn. Feedback particle filter. IEEE Transactions on Automatic Control. vol. 58, no. 10. pp. 2465–2480 (2013).

[43] T. Yang, R.S. Laugesen, P.G. Mehta, and S.P. Meyn. Multivariable feedback particle filter. Automatica. vol. 71. pp. 10–23 (2016).

[44] M. Zakai. On the Optimal Filtering of Diffusion Processes. Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete. vol. 11, no. 3. pp. 230–243 (1969).

b

THE EFFECTIVE SAMPLE SIZE CRITERION

The effective sample size criterion [49] measures the loss of efficiency using the importance sampling estimator $\tilde{\mu}(\boldsymbol{X})$ (3.7) instead of the Monte Carlo estimator $\bar{\mu}(\boldsymbol{X})$ (3.2) defined in Section 3.1:

$$\bar{\mu}(\boldsymbol{X}) = \frac{1}{N}\sum_{i=1}^{N} \mu(\boldsymbol{X}^i), \quad \boldsymbol{X}^i \sim p \tag{b.1}$$

$$\tilde{\mu}(\boldsymbol{X}) = \sum_{i=1}^{N} w^i \, \mu(\boldsymbol{X}^i), \quad \boldsymbol{X}^i \sim q \tag{b.2}$$

where $(\boldsymbol{X}^i)_{i=1,\dots,N}$ is a set of independent random variables on $\mathbb{R}^d$ with the same distribution as $\boldsymbol{X}$, q is the proposal density, $\tilde{w}^i = \frac{p(\boldsymbol{X}^i)}{q(\boldsymbol{X}^i)}$ are the unnormalised weight and $w^i = \frac{\tilde{w}^i}{\sum_{j=1}^{N} \tilde{w}^j}$ the normalised one, and $\mu : \mathbb{R}^d \to \mathbb{R}$ is a bounded function.

The effective sample size is defined by:

$$ESS = N \, \frac{\mathbb{V}\mathrm{ar}_p(\bar{\mu}(\boldsymbol{X}))}{\mathbb{V}\mathrm{ar}_q(\tilde{\mu}(\boldsymbol{X}))} \tag{b.3}$$

However, in the literature, the following expression is more commonly found:

$$\mathrm{ESS} = \frac{N}{1 + \mathbb{V}\mathrm{ar}_q(\tilde{w})} \tag{b.4}$$

and the approximation of the criterion is often given by:

$$\widehat{ESS} = \frac{1}{\sum_{i=1}^{N} (w^i)^2} \tag{b.5}$$

Section b.1 provides the elements to pass from equation (b.3) to (b.4). The passage from the effective sample size criterion equation (b.4) to the approximate criterion (b.5) is provided in Section b.2.

B.1 EXPRESSIONS OF THE EFFECTIVE SAMPLE SIZE CRITERION

We start by approximating the variance of the importance sampling estimator $\mathbb{V}\mathrm{ar}_q(\tilde{\mu}(\boldsymbol{X}))$. Using the definition of the normalised weights, it comes:

$$\mathbb{V}\mathrm{ar}_q(\tilde{\mu}(\boldsymbol{X})) = \mathbb{V}\mathrm{ar}_q\left(\frac{\sum_{i=1}^N \tilde{w}^i \mu(\boldsymbol{X}^i)}{\sum_{i=1}^N \tilde{w}^i}\right) \tag{b.6}$$

The variance of the importance sampling estimator can be approximated using the delta-method [8, 66] (a.k.a., Taylor expansions for the moments of functions of random variables):

$$\begin{aligned}\mathbb{V}\mathrm{ar}_q\left(\frac{\sum_{i=1}^N \tilde{w}^i \mu(\boldsymbol{X}^i)}{\sum_{i=1}^N \tilde{w}^i}\right) \approx \frac{1}{N\,(\mathbb{E}_q(\tilde{w}))^2}\Big(&\mathbb{V}\mathrm{ar}_q(\tilde{w}\,\mu(\boldsymbol{X})) \\ &-2\,\frac{\mathbb{E}_q(\tilde{w}\,\mu(\boldsymbol{X}))}{\mathbb{E}_q(\tilde{w})}\,\mathrm{COV}_q(\tilde{w}, \tilde{w}\,\mu(\boldsymbol{X})) + \left(\frac{\mathbb{E}_q(\tilde{w}\,\mu(\boldsymbol{X}))}{\mathbb{E}_q(\tilde{w})}\right)^2 \mathbb{V}\mathrm{ar}_q(\tilde{w})\Big)\end{aligned} \tag{b.7}$$

We pose $\mathbb{E}_q(\tilde{w}) = 1$. By taking $\boldsymbol{I} \triangleq \mathbb{E}_p(\mu(\boldsymbol{X}))$ and noticing that

$$\begin{aligned}\mathbb{E}_q(\tilde{w}\,\mu(\boldsymbol{X})) &= \int \tilde{w}\,\mu(\boldsymbol{X})\,q(\boldsymbol{X})\,d\boldsymbol{X} && \text{(b.8)}\\ &= \int \frac{p(\boldsymbol{X})}{q(\boldsymbol{X})}\,\mu(\boldsymbol{X})\,q(\boldsymbol{X})\,d\boldsymbol{X} && \text{(b.9)}\\ &= \mathbb{E}_p(\mu(\boldsymbol{X})) = \boldsymbol{I} && \text{(b.10)}\end{aligned}$$

the variance becomes:

$$\begin{aligned}\mathbb{V}\mathrm{ar}_q\left(\frac{\sum_{i=1}^N \tilde{w}^i\,\mu(\boldsymbol{X}^i)}{\sum_{i=1}^N \tilde{w}^i}\right) \approx \frac{1}{N}\,\big(&\mathbb{V}\mathrm{ar}_q(\tilde{w}\,\mu(\boldsymbol{X})) \\ &-2\,\boldsymbol{I}\,\mathrm{COV}_q(\tilde{w}, \tilde{w}\,\mu(\boldsymbol{X})) + \boldsymbol{I}^2\,\mathbb{V}\mathrm{ar}_q(\tilde{w})\big)\end{aligned} \tag{b.11}$$

Studying the terms separately, it comes:

$$\begin{aligned}\mathrm{COV}_q(\tilde{w}, \tilde{w}\,\mu(\boldsymbol{X})) &= \mathbb{E}_q(\tilde{w}^2\,\mu(\boldsymbol{X})) - \boldsymbol{I} && \text{(b.12)}\\ &= \mathbb{E}_p(\tilde{w}\,\mu(\boldsymbol{X})) - \boldsymbol{I} && \text{(b.13)}\\ &= \mathbb{E}_p(\tilde{w}\,\mu(\boldsymbol{X})) - \mathbb{E}_p(\tilde{w})\,\mathbb{E}_p(\mu(\boldsymbol{X})) && \text{(b.14)}\\ &\quad + \mathbb{E}_p(\tilde{w})\,\mathbb{E}_p(\mu(\boldsymbol{X})) - \boldsymbol{I} \\ &= \mathrm{COV}_p(\tilde{w}, \mu(\boldsymbol{X})) + \boldsymbol{I}\,\mathbb{E}_p(\tilde{w}) - \boldsymbol{I} && \text{(b.15)}\end{aligned}$$

Equation (b.13) comes from:

$$\begin{aligned}\mathbb{E}_q(\tilde{w}^2\,\mu(\boldsymbol{X})) &= \int \tilde{w}^2\,\mu(\boldsymbol{X})q(\boldsymbol{X})\,d\boldsymbol{X} && \text{(b.16)}\\ &= \int \tilde{w}\,\mu(\boldsymbol{X})p(\boldsymbol{X})\,d\boldsymbol{X} && \text{(b.17)}\\ &= \mathbb{E}_p(\tilde{w}\,\mu(\boldsymbol{X})) && \text{(b.18)}\end{aligned}$$

Another term of the variance (b.11) is $\mathbb{V}\mathrm{ar}_q(\tilde{w}\,\mu(\boldsymbol{X}))$ that can be rewitten as:

$$\mathbb{V}\mathrm{ar}_q(\tilde{w}\,\mu(\boldsymbol{X})) = \mathbb{E}_q((\tilde{w}\,\mu(\boldsymbol{X}))^2) - \boldsymbol{I}^2 \tag{b.19}$$
$$= \mathbb{E}_p(\tilde{w}\,\mu(\boldsymbol{X})^2) - \boldsymbol{I}^2 \tag{b.20}$$

Indeed, we have:

$$\mathbb{E}_q((\tilde{w}\,\mu(\boldsymbol{X}))^2) = \int \mu(\boldsymbol{X})^2\,\tilde{w}^2\,q(\boldsymbol{X})\,d\boldsymbol{X} \tag{b.21}$$
$$= \int \mu(\boldsymbol{X})^2\,\tilde{w}\,\frac{p(\boldsymbol{X})}{q(\boldsymbol{X})}\,q(\boldsymbol{X})\,d\boldsymbol{X} \tag{b.22}$$
$$= \mathbb{E}_p(\tilde{w}\,\mu(\boldsymbol{X})^2) \tag{b.23}$$

The term $\mathbb{E}_p(\tilde{w}\,\mu(\boldsymbol{X})^2)$ can be approximating using again the delta-method:

$$\mathbb{E}_p(\tilde{w}\,\mu(\boldsymbol{X})^2) \approx \mathbb{E}_p(\tilde{w})\,\boldsymbol{I}^2 + \mathbb{E}_p(\tilde{w})\,\mathbb{V}\mathrm{ar}_p(\mu(\boldsymbol{X})) + 2\,\boldsymbol{I}\,\mathrm{COV}_p(\tilde{w}), \mu(\boldsymbol{X})) \tag{b.24}$$

Therefore, by plugging equations (b.15), (b.20) and (b.24) in the expression of the variance (b.11), it comes:

$$\mathbb{V}\mathrm{ar}_q\left(\frac{\sum_{i=1}^N \tilde{w}^i \mu(\boldsymbol{X}^i)}{\sum_{i=1}^N \tilde{w}^i}\right) \approx \frac{1}{N}\left(\mathbb{V}\mathrm{ar}_p(\mu(\boldsymbol{X}))\,\mathbb{E}_p(\tilde{w}) + \boldsymbol{I}^2\,(1 + \mathbb{V}\mathrm{ar}_q(\tilde{w}) - \mathbb{E}_p(\tilde{w}))\right) \tag{b.25}$$

We have:

$$\mathbb{V}\mathrm{ar}_q(\tilde{w}) = \mathbb{E}_q(\tilde{w}^2) - 1 \tag{b.26}$$

and

$$\mathbb{E}_q(\tilde{w}^2) = \int \tilde{w}^2 q(\boldsymbol{X})\,d\boldsymbol{X} = \int \tilde{w} p(\boldsymbol{X})\,d\boldsymbol{X} = \mathbb{E}_p(\tilde{w}) \tag{b.27}$$

Hence, the approximation of the variance of the importance sampling estimator:

$$\mathbb{V}\mathrm{ar}_q(\tilde{\mu}(\boldsymbol{X})) \approx \frac{1}{N}\,\mathbb{V}\mathrm{ar}_p(\mu(\boldsymbol{X}))\,(1 + \mathbb{V}\mathrm{ar}_q(\tilde{w})) \tag{b.28}$$

By taking this approximation (b.28), the effective sample size criterion is approximated by:

$$ESS = N\,\frac{\mathbb{V}\mathrm{ar}_p(\bar{\mu}(\boldsymbol{X}))}{\mathbb{V}\mathrm{ar}_q(\tilde{\mu}(\boldsymbol{X}))} \tag{b.29}$$
$$= N\,\frac{\frac{1}{N}\mathbb{V}\mathrm{ar}_p(\mu(\boldsymbol{X}))}{\frac{1}{N}\mathbb{V}\mathrm{ar}_p(\mu(\boldsymbol{X}))\,(1 + \mathbb{V}\mathrm{ar}_q(\tilde{w}))} \tag{b.30}$$
$$= \frac{N}{1 + \mathbb{V}\mathrm{ar}_q(\tilde{w})} \tag{b.31}$$

B.2 APPROXIMATION OF THE EFFECTIVE SAMPLE SIZE CRITERION

The effective sample size (b.4) cannot be evaluated exactly. The variance of the unnormalised weights is first replaced by the coefficient of variance:

$$\mathrm{ESS} \approx \frac{N}{1 + \frac{\mathbb{V}\mathrm{ar}_q(\tilde{w})}{\mathbb{E}_q(\tilde{w})^2}} \tag{b.32}$$

By definition of the variance:

$$1 + \frac{\mathbb{V}\mathrm{ar}_q(\tilde{w})}{\mathbb{E}_q(\tilde{w})^2} = 1 + \frac{\mathbb{E}_q(\tilde{w}^2) - \mathbb{E}_q(\tilde{w})^2}{\mathbb{E}_q(\tilde{w})^2} = \frac{\mathbb{E}_q(\tilde{w}^2)}{\mathbb{E}_q(\tilde{w})^2} \tag{b.33}$$

Approximating expectations by sums, it follows:

$$\mathrm{ESS} = \frac{N}{\frac{\mathbb{E}_q(\tilde{w}^2)}{\mathbb{E}_q(\tilde{w})^2}} \simeq \frac{N}{\frac{\frac{1}{N}\sum_{i=1}^N (\tilde{w}^i)^2}{\left(\frac{1}{N}\sum_{i=1}^N \tilde{w}^i\right)^2}} \tag{b.34}$$

The estimate of the effective sample size [49] is thus given by:

$$\widehat{ESS} = \frac{1}{\sum_{i=1}^N (w^i)^2} \tag{b.35}$$

where $w^i = \frac{\tilde{w}^i}{\sum_{i=1}^N \tilde{w}^i}$ are the normalised weights.

C

INTRODUCTION TO INFORMATION THEORY

The Fisher information matrix aims to mathematically quantify the informativeness carried by a random vector (e.g., a measurement or an estimate). Information is usually derived from the associated probability density function ant quantifies its compactness [36].

C.1 DEFINITION OF THE FISHER INFORMATION MATRIX

Let $p(\boldsymbol{y}|\boldsymbol{\theta})$ be the conditional density for a random vector $\boldsymbol{y}$ taking values in $\mathbb{R}^{d_y}$ conditioned on the value of $\boldsymbol{\theta}$. We assume that the two derivatives of the conditional density $p(\boldsymbol{y}|\boldsymbol{\theta})$ exist, and that the support of the density is independent from $\boldsymbol{\theta}$. Under this regularity conditions, the Fisher information matrix is defined by the variance of the score function [52]:

$$\mathcal{I}(\boldsymbol{\theta}) = -\mathbb{E}_{y|\theta}\left[\frac{\partial^2}{\partial \boldsymbol{\theta}^2}\log\left(p(\boldsymbol{y}|\boldsymbol{\theta})\right)\right] \tag{c.1}$$

where $\mathbb{E}_{y|\theta}[f(\boldsymbol{y},\boldsymbol{\theta})] = \int_{\mathbb{R}^{d_y}} f(\boldsymbol{y},\boldsymbol{\theta})\, p(\boldsymbol{y}|\boldsymbol{\theta})\, d\boldsymbol{y}$ is the mathematical expectation. The Fisher information matrix can be interpreted as the expectation of the likelihood curvature with respect to $\boldsymbol{\theta}$.

The Fisher information matrix is used to quantify the information brought by the measurements to the state in the context of state estimation. In this case, the conditional density is the likelihood $p(\boldsymbol{y}|\boldsymbol{x})$, where $\boldsymbol{y} \in \mathbb{R}^{d_y}$ is the measurement, and $\boldsymbol{x} \in \mathbb{R}^d$ is the state. The inverse of the Fisher information matrix is usually referred to as the PCRB and quantifies the lowest covariance reachable by any unbiased estimator for a given estimation problem.

C.2 POSTERIOR CRAMÉR-RAO BOUND

The PCRB vectorial form in the discrete-case has been introduced by Galdos [30]. In order to reduce the computational cost, an iterative formulation for recursive estimation has been derived by Tichavsky, Muravchik, and Nehorai [84]. Let the state-space model be defined by:

$$\begin{cases} \boldsymbol{x}_{k+1} = f_k(\boldsymbol{x}_k) + \boldsymbol{\eta}_k \\ \boldsymbol{y}_k = h_k(\boldsymbol{x}_k) + \boldsymbol{\nu}_k \end{cases} \tag{c.2}$$

where the noises $\boldsymbol{\eta}_k$ and $\boldsymbol{\nu}_k$ are noises with covariance $\boldsymbol{Q}_k$ and $\boldsymbol{R}_k$ respectively.

At each time-step, the PCRB is given by the following Ricatti-like recursion:

$$PCRB_{k+1} = \left(\mathcal{I}(\boldsymbol{x}_{k+1})\right)^{-1} \tag{c.3a}$$

$$\mathcal{I}(\boldsymbol{x}_{k+1}) = \boldsymbol{D}_k^{22} - \boldsymbol{D}_k^{12\top}\left(\mathcal{I}(\boldsymbol{x}_k) + \boldsymbol{D}_k^{11}\right)^{-1}\boldsymbol{D}_k^{12} \tag{c.3b}$$

where the matrices $\boldsymbol{D}_k^{11}$, $\boldsymbol{D}_k^{12}$, and $\boldsymbol{D}_k^{22}$ are defined by:

$$\boldsymbol{D}_k^{11} = \mathbb{E}\left[-\frac{\partial^2}{\partial \boldsymbol{x}_k^2}\log p(\boldsymbol{x}_{k+1}|\boldsymbol{x}_k)\right] \tag{c.4a}$$

$$\boldsymbol{D}_k^{12} = \mathbb{E}\left[-\frac{\partial^2}{\partial \boldsymbol{x}_k \boldsymbol{x}_{k+1}}p(\boldsymbol{x}_{k+1}|\boldsymbol{x}_k)\right] \tag{c.4b}$$

$$\boldsymbol{D}_k^{22} = \mathbb{E}\left[-\frac{\partial^2}{\partial \boldsymbol{x}_{k+1}^2}p(\boldsymbol{x}_{k+1}|\boldsymbol{x}_k)\right] + \mathbb{E}\left[-\frac{\partial^2}{\partial \boldsymbol{x}_{k+1}^2}p(\boldsymbol{y}_{k+1}|\boldsymbol{x}_{k+1})\right] \tag{c.4c}$$

When the state model is linear, i.e., $\boldsymbol{x}_{k+1} = \boldsymbol{F}_k\,\boldsymbol{x}_k + \boldsymbol{\eta}_k$ and the noises Gaussian, the PCRB formulation (c.3b) is simplified by [9]:

$$\mathcal{I}(\boldsymbol{x}_{k+1}) = \mathbb{E}\left[\left[\frac{\partial h_{k+1}^\top(\boldsymbol{x}_{k+1})}{\partial \boldsymbol{x}_{k+1}}\right]\boldsymbol{R}_{k+1}^{-1}\left[\frac{\partial h_{k+1}^\top(\boldsymbol{x}_{k+1})}{\partial \boldsymbol{x}_{k+1}}\right]^\top\right] + \left(\boldsymbol{Q}_k + \boldsymbol{F}_k\,\mathcal{I}(\boldsymbol{x}_k)^{-1}\boldsymbol{F}_k^\top\right)^{-1} \tag{c.5}$$

The initial Fisher information matrix $\mathcal{I}(\boldsymbol{x}_0)$ is equal to the covariance $\boldsymbol{P}_0$ of the initial state density.

As the measurements do not appear in the computation of the PCRB, the bound can be calculated offline. In practice, as the true state $\boldsymbol{x}_k$ is unknown, the PCRB is approximated. The expectation terms are computed through Monte Carlo simulations of the state.

d

GRAVIMETER OBSERVATION EQUATION FOR THE DOUBLE INTEGRATOR MODEL

The observation equation for the double integrator model is derived from the equation (4.31):

$$Y^{g'} = g_{nom}(p_\lambda, p_h) \; + \; map_{ga}(p_\phi, p_\lambda) \; + \; \ddot{p}_h \; + \; Eötvös(V_E, p_\lambda) \; + \; \nu \tag{d.1}$$

In order to transposed the observation equation (d.1) in the double integrator model settings, the nominal gravity and Eötvös effect equations (4.29) and (4.30) are linearised by using the first order Taylor's theorem:

$$\begin{aligned} g_{nom}(p_\lambda, p_h) = g_{nom}(p_{\lambda 0}, p_{h0}) + 1.0326 \times 10^4 \sin(p_{\lambda 0}) \cos(p_{\lambda 0})(p_\lambda - p_{\lambda 0}) \\ + 0.3086 \; (p_h - p_{h0}) + O(||(p_\lambda - p_{\lambda 0}, \, p_h - p_{h0})||_2) \end{aligned} \tag{d.2a}$$

$$\begin{aligned} Eötvös(V_E, p_\lambda) = Eötvös(V_{E_0}, \lambda_0) + 3.7515 \cos(\lambda_0)(V_E - V_{E_0}) \\ - 3.7515 \; V_{E_0} \; \sin(\lambda_0)(\lambda - \lambda_0) + O(||(V_E - V_{E_0}, \, \lambda - \lambda_0)||_2) \end{aligned} \tag{d.2b}$$

The latitude $p_{\lambda 0} = 0.6196$ rad ($p_{\lambda 0} = 35.5°$ N) and the longitude $p_{\phi 0} = -2.1642$ rad ($p_{\phi 0} = 124°$ W) correspond to the reference point of the maps of the Californian coasts (see Figure 4.7). To facilitate the calculations, the velocity V_{E_0} and the altitude p_{h0} are taken equal to 0. By simplifying,

$$\begin{aligned} g_{nom}(p_\lambda, p_h) = g_{nom}(p_{\lambda 0}, p_{h0}) + 4.8816 \times 10^3 (p_\lambda - p_{\lambda 0}) + 0.3086 \; (p_h - p_{h0}) \\ + O(||(p_\lambda - p_{\lambda 0}, p_h - p_{h0})||_2) \end{aligned} \tag{d.3a}$$

$$\begin{aligned} Eötvös(V_E, p_\lambda) = \text{Eötvös}(V_{E_0}, p_{\lambda 0}) + 3.0542 \; (V_E - V_{E_0}) \\ + O(||(V_E - V_{E_0}, p_\lambda - p_{\lambda 0})||_2) \end{aligned} \tag{d.3b}$$

The equations are expressed with coordinates on Flat Earth using the following relations:

$$\begin{cases} (p_x - p_{x0}) = (R_\phi + p_{h0}) \; \cos(p_{\lambda 0}) \; (p_\phi - p_{\phi 0}) \\ (p_y - p_{y0}) = (R_\lambda + p_{h0}) \; (p_\lambda - p_{\lambda 0}) \\ (p_z - p_{z0}) = -(p_h - p_{h0}) \end{cases} \tag{d.4}$$

where R_λ is the radius of curvature of the Earth in the meridian plane and R_ϕ is the large normal of the ellipsoid. In the following R_λ and R_ϕ are approximated by $R = 6371 \times 10^3$ m which is the radius of the Earth. The chosen reference point $[p_{x0}, p_{y0}, p_{z0}]^\top$ is $[0, 0, 0]^\top$. The east velocity V_E corresponds to the velocity in the x-axis V_x. Using these relations (d.4), equations (d.3a) and (d.3a) become:

$$g_{nom}(p_\lambda, p_h) = g_{nom}(p_{\lambda 0}, p_{h0}) + \frac{4.8816 \times 10^3}{R} p_y - 0.3086\ p_z + O(||(p_y, p_z)||_2) \tag{d.5a}$$

$$\textit{Eötvös}(V_E, p_\lambda) = \text{Eötvös}(V_{E_0}, p_{\lambda 0}) + 3.0542\ V_x + O(||(V_x, p_y)||_2) \tag{d.5b}$$

Let Y^g be a new observation equation such that:

$$Y^g = Y^{g'} - g_{nom}(p_{\lambda 0}, p_{h0}) - \text{Eötvös}(V_{E_0}, p_{\lambda 0}) \tag{d.6}$$

The gravimeter observation equation for the double integrator model (4.32) comes from taking into account the simplifications of the mathematical models of the nominal gravity and Eötvös effect.

e

FISHER INFORMATION MATRIX APPLIED TO UNDERWATER NAVIGATION

In Chapter 4, the observability level of the underwater navigation system is quantified by the Fisher information matrix (see Section 4.2.3, equation (4.33)). The measurement contribution to the Fisher information matrix is provided for the multi-beam telemeter (4.36), and the multi-beam telemeter and gravimeter (4.38).

In this appendix, the full Fisher information matrix are given under the following assumptions:

- The bathymetry equation is also approaches as a single-beam telemeter pointing towards the local vertical direction (4.35):

$$Y^b = p_z - \text{map}_{\text{mb}}(p_x, p_y) + \nu \tag{e.1}$$

- The process noise covariance $\boldsymbol{Q}_k$ is assumed to be equal to $\mathbf{0}_6$.
- The inverse of the Fisher information matrix at time k, $\mathcal{I}(\boldsymbol{X}_k)^{-1}$ is equal to the inverse of the initial uncertainty matrix $\boldsymbol{P}_0^{-1}$:

$$\boldsymbol{P}_0^{-1} = \begin{pmatrix} \sigma_p^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_p^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_p^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_V^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_V^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_V^2 \end{pmatrix} \tag{e.2}$$

  To simplify, the standard derivation of the position and velocity is the same for all axis (x, y and z).

The Fisher information matrices are calculated in order to compare their determinant. If the determinant of the information matrix when the gravimeter is activated is greater than the determinant of the information matrix with only the multi-beam telemeter, then the additional sensor brings more information and thus leads to more accurate estimates.

Under the above assumptions, the dynamic contribution to information is given by:

$$\left(\boldsymbol{B}_k\boldsymbol{P}_0^{-1}\boldsymbol{B}_k^\top\right)^{-1} = \begin{bmatrix} \sigma_p^2+\sigma_v^2dt^2 & 0 & 0 & \sigma_v^2dt & 0 & 0 \\ 0 & \sigma_p^2+\sigma_v^2dt^2 & 0 & 0 & \sigma_v^2dt & 0 \\ 0 & 0 & \sigma_p^2+\sigma_v^2dt^2 & 0 & 0 & \sigma_v^2dt \\ \sigma_v^2dt & 0 & 0 & \sigma_v^2 & 0 & 0 \\ 0 & \sigma_v^2dt & 0 & 0 & \sigma_v^2 & 0 \\ 0 & 0 & \sigma_v^2dt & 0 & 0 & \sigma_v^2 \end{bmatrix} \tag{e.3}$$

Let the constant d be equal to $-\sigma_v^4dt^2+\sigma_v^2\left(\sigma_p^2+\sigma_v^2dt^2\right)$. The Fisher information matrix for the bathymetric model, is:

$$J_{\text{bathy}} = \begin{bmatrix} \frac{\left(\nabla_b^x\right)^2}{\sigma_b^2}+\frac{\sigma_v^2}{d} & \frac{\nabla_b^x\nabla_b^y}{\sigma_b^2} & -\frac{\nabla_b^x}{\sigma_b^2} & -\frac{\sigma_v^2\,dt}{d} & 0 & 0 \\ \frac{\nabla_b^x\nabla_b^y}{\sigma_b^2} & \frac{\left(\nabla_b^y\right)^2}{\sigma_b^2}+\frac{\sigma_v^2}{d} & -\frac{\nabla_b^y}{\sigma_b^2} & 0 & -\frac{\sigma_v^2\,dt}{d} & 0 \\ -\frac{\nabla_b^x}{\sigma_b^2} & -\frac{\nabla_b^y}{\sigma_b^2} & \frac{1}{\sigma_b^2}+\frac{\sigma_v^2}{d} & 0 & 0 & -\frac{\sigma_v^2\,dt}{d} \\ -\frac{\sigma_v^2\,dt}{d} & 0 & 0 & \frac{\sigma_p^2+\sigma_v^2dt^2}{d} & 0 & 0 \\ 0 & -\frac{\sigma_v^2dt}{d} & 0 & 0 & \frac{\sigma_p^2+\sigma_v^2dt^2}{d} & 0 \\ 0 & 0 & -\frac{\sigma_v^2dt}{d} & 0 & 0 & \frac{\sigma_p^2+\sigma_v^2dt^2}{d} \end{bmatrix} \tag{e.4}$$

where ∇_{mb}^x and ∇_{mb}^y are the East and North gradients of the seabed elevation map, and σ_{mb} is the standard deviation of the telemeter measurement error.

The Fisher information matrix when the gravity field measurement is added, is given by:

$$J_{\text{fusion}} = \begin{bmatrix} \frac{\left(\nabla_b^x\right)^2}{\sigma_b^2}+\frac{\left(\nabla_g^x\right)^2}{\sigma_g^2}+\frac{\sigma_v^2}{d} & \frac{\nabla_b^x\nabla_b^y}{\sigma_b^2}+\frac{\nabla_g^x(\nabla_g^y+a)}{\sigma_g^2} & -\frac{\nabla_b^x}{\sigma_b^2}+\frac{\nabla_g^xb}{\sigma_g^2} \\ \frac{\nabla_b^x\nabla_b^y}{\sigma_b^2}+\frac{\nabla_g^x(\nabla_g^y+a)}{\sigma_g^2} & \frac{\left(\nabla_b^y\right)^2}{\sigma_b^2}+\frac{\sigma_v^2}{d}+\frac{\left(\nabla_g^y+a\right)^2}{\sigma_g^2} & -\frac{\nabla_b^y}{\sigma_b^2}+\frac{b(\nabla_g^y+a)}{\sigma_g^2} \\ -\frac{\nabla_b^x}{\sigma_b^2}+\frac{\nabla_g^xb}{\sigma_g^2} & -\frac{\nabla_b^y}{\sigma_b^2}+\frac{b(\nabla_g^y+a)}{\sigma_g^2} & \frac{\sigma_v^2}{d}+\frac{b^2}{\sigma_g^2}+\frac{1}{\sigma_b^2} \\ \frac{\nabla_g^xc}{\sigma_g^2}-\frac{\sigma_v^2dt}{d} & \frac{c(\nabla_g^y+a)}{\sigma_g^2} & \frac{bc}{\sigma_g^2} \\ 0 & -\frac{\sigma_v^2dt}{d} & 0 \\ 0 & 0 & -\frac{\sigma_v^2dt}{d} \end{bmatrix}$$

$$\begin{bmatrix} \frac{\nabla_g^xc}{\sigma_g^2}-\frac{\sigma_v^2dt}{d} & 0 & 0 \\ \frac{c(\nabla_g^y+a)}{\sigma_g^2} & -\frac{\sigma_v^2dt}{d} & 0 \\ \frac{bc}{\sigma_g^2} & 0 & -\frac{\sigma_v^2dt}{d} \\ \frac{\sigma_p^2+\sigma_v^2dt^2}{d}+\frac{c^2}{\sigma_g^2} & 0 & 0 \\ 0 & \frac{\sigma_p^2+\sigma_v^2dt^2}{d} & 0 \\ 0 & 0 & \frac{\sigma_p^2+\sigma_v^2dt^2}{d} \end{bmatrix} \tag{e.5}$$

where ∇^x_{ga} and ∇^y_{ga} are the East and North gradients of the local gravimetry anomalies map, and σ_{ga} is the standard deviation of the atomic gravimeter measurement error.

Determinants of both matrices are given in Equations (4.39) and (4.40), and are recalled here:

$$\det(\mathcal{I}^b) = \frac{\alpha_{mb}\,\Delta t^2 + \beta_{mb}}{q_{mb}} \tag{e.6}$$

$$\det(\mathcal{I}^f) = \frac{\alpha_{ga}\,\Delta t^2 + \beta_{ga}\Delta t + \gamma_{ga} + o(\Delta_t^2, {\nabla^x_{mb}}^2, {\nabla^x_{ga}}^2)}{q_{ga}} \tag{e.7}$$

where the constants are:

$$\alpha_{mb} = {\nabla^x_{mb}}^2\sigma_V + {\nabla^y_{mb}}^2\sigma_V + \sigma_V \tag{e.8}$$

$$\beta_{mb} = {\nabla^x_{mb}}^2\sigma_p + {\nabla^y_{mb}}^2\sigma_p + \sigma^2_{mb} + \sigma_p \tag{e.9}$$

$$q_{mb} = \sigma^2_{mb}\,\sigma^3_p\,\sigma^3_V \tag{e.10}$$

$$\begin{aligned}\alpha_g = 2\sigma^2_p\sigma^2_V \Big(&{\nabla^x_{mb}}^2a^2 + {\nabla^x_{mb}}^2b^2 + 2\nabla^x_{mb}\nabla^x_{ga}b + {\nabla^y_{mb}}^2b^2 + 2\nabla^y_{mb}\nabla^y_{ga}b + 2\nabla^y_{mb}ab \\ &+{\nabla^x_{ga}}^2 + {\nabla^y_{ga}}^2 + 2\nabla^y_{ga}a + a^2\Big) + \sigma^2_{ga}\sigma^2_V\left({\nabla^x_{mb}}^2 + {\nabla^y_{mb}}^2 + 1\right) \\ &+ \sigma^4_V c^2\left({\nabla^y_{mb}}^2 + 1\right) + \sigma^2_{mb}\sigma^2_V\left({\nabla^x_{ga}}^2 + {\nabla^y_{ga}}^2 + 2\nabla^y_{ga}a + a^2 + b^2\right)\end{aligned} \tag{e.11}$$

$$\beta_g = \sigma^2_p\sigma^2_V c\left(-2\nabla^x_{mb}\nabla^y_{mb}a + 2\nabla^x_{mb}b + 2\nabla^x_{ga}\right) + 2\nabla^x_{ga}\sigma^2_{mb}\sigma^2_V c \tag{e.12}$$

$$\begin{aligned}\gamma_g = \sigma^2_{ga}\sigma^2_p\left({\nabla^x_{mb}}^2 + {\nabla^y_{mb}}^2 + 1\right) + \sigma^4_p\Big(&{\nabla^x_{mb}}^2a^2 + {\nabla^x_{mb}}^2b^2 + 2\nabla^x_{mb}\nabla^x_{ga}b + {\nabla^y_{mb}}^2b^2 \\ &+2\nabla^y_{mb}\nabla^y_{ga}b + 2\nabla^y_{mb}ab + {\nabla^x_{ga}}^2 + {\nabla^y_{ga}}^2 + 2\nabla^y_{ga}a + a^2\Big) \\ &+ \sigma^2_p\sigma^2_V\left({\nabla^x_{mb}}^2c^2 + {\nabla^y_{mb}}^2c^2 + c^2\right) + \sigma^2_{mb}\sigma^2_p\left({\nabla^x_{ga}}^2 + {\nabla^y_{ga}}^2 + 2\nabla^y_{ga}a\right. \\ &\left.+a^2 + b^2\right) + \sigma^2_{mb}\sigma^2_{ga} + \sigma^2_{mb}\sigma^2_V c^2\end{aligned} \tag{e.13}$$

$$q_g = \sigma^2_{mb}\sigma^2_{ga}\sigma^6_p\sigma^6_V \tag{e.14}$$

Any quadratic product of the gradient of the maps of order greater than or equal to 3 were removed of the determinant in equation (4.40).

f

COMPLEXITY ANALYSIS

The complexity of algorithms can be evaluated in terms of their computational load for a given set of numerical operations. In this work, the computational load is defined as the total number of floating point operations (or flops, namely additions and multiplications) required to perform them during one time-step (i.e., prediction, correction and resampling). The computational cost of the initialization can be neglected since this step is done only once. We consider the most expensive time-step in term of flops as the cost of the regularised resampling is included although it is not triggered at each time-step. This complexity criterion is often used in the state estimation area [46]. Table f.1 presents the flops number for the elementary operations used in this paper.

Table f.1: The flops complexity associated with elementary operations.

Instruction	Size	Multiplications	Additions
$A + B$	$A, B \in \mathbb{R}^{n\times m}$		nm
AB	$A \in \mathbb{R}^{n\times m}$, $B \in \mathbb{R}^{m\times l}$	nml	$(m-1)nl$
A^{-1}	$A \in \mathbb{R}^{n\times n}$	n^3	
$chol(A)$	$A \in \mathbb{R}^{n\times n}$	$n^3/3 + 2n^2$	

Some operations cannot be quantified in terms of flops. Thus, the computational load of nonlinear measurements model is noted c_g, the complexity of the resampling and the regularisation $c_{\text{reg-res}}$ and the theoretical cost of one random sample is denoted c_{random} without distinction of distributions. The determination of the A2BC bandwidth parameter ε is denoted c_ε. The computational load of the proposal density is denoted c_q and the one of the prior density is c_p.

The computation for the RBPF yields the following number of flops per time-step, assuming a linear dynamics [46]:

$$c_{RBPF} = N \Big(6nl + 4n^2 + 2l^2 + n - l + n\, c_{\text{random}} + c_{\text{reg-res}} + c_g + 4nl^2 + 8ln^2 + 4/3n^3 + 5l^3 - 5nl + 2n^2 + l^3\Big) \quad \text{(f.1)}$$

where n and l are the dimension of respectively the nonlinear and the linear parts of the state vector. The computational load of the RPF can easily be derived from the RBPF under the same assumptions [46]:

$$c_{RPF} = N \left(6d^2 - d + dc_{\text{random}} + c_{\text{reg-res}} + c_g + 3\right) + \frac{d^3}{3} + 2d^2 - 1 \tag{f.2}$$

The number of particles is the most important factor in terms of computational load. The RBPF needs less particles than the RPF as the dimension of the nonlinear state vector is small, implying that the particles occupy a lower dimensional space. Given a computational complexity, the number of particles needed by a standard particle filter depending on the number of particles of the RBPF (without the regularisation step) is given in [46].

g

ADAPTIVE APPROXIMATE BAYESIAN COMPUTATION PARTICLE FILTERS

Algorithm g.1 A2BC-RPF

Initialisation: The initial particle set is drawn as $\{\boldsymbol{X}_0^i\}_{i\in[1,N]}$ using the initial density $p(\boldsymbol{X}_0)$ and the initial weights set $\{w_0^i\}_{i\in[1,N]}$ is taken equal to $\frac{1}{N}$.
for each time-step k **do**
  Prediction: Draw the particles from the transition density $\boldsymbol{X}_k^i \sim p(\boldsymbol{X}_k|\boldsymbol{X}_{k-1}^i)$.
  A2BC correction: Determine ϵ_k. Update the weights $w_k^i \propto w_{k-1}^i \mathcal{K}_{\epsilon_k}(\boldsymbol{Y}_k - h_k(\boldsymbol{X}_k^i))$.
  Compute the state estimate $\hat{\boldsymbol{X}}_k$ (3.16) and its covariance $\boldsymbol{P}_k$ (3.17).
  if a resampling criterion is satisfied, e.g., $\widehat{\mathrm{ESS}} < N_{\mathrm{th}}$, see (3.19) **then**
    Draw a new set of particles $\{\boldsymbol{X}_k^i\}_{i\in[1,N]}$ using a resampling method, e.g., the multinomial resampling (Algorithm 3.2).
    Reset the weights to $\frac{1}{N}$.
  end if
end for
Return the state estimate $\hat{\boldsymbol{X}}_k$ and its covariance $\boldsymbol{P}_k$, $\forall k$.

Algorithm g.2 A2BC-RBPF

Initialisation: The initial particles set is drawn as $\{\boldsymbol{X}_0^{n,i}\}_{i\in[1,N]}$ using the initial density $p(\boldsymbol{X}_0^n)$ and the initial weights set $\{w_0^i\}_{i\in[1,N]}$ is taken equal to $\frac{1}{N}$. The initial mean and covariance of the linear state variables set is $\left\{\boldsymbol{X}_{0|0}^{l,i}, \boldsymbol{P}_{0|0}^{l,i}\right\} = \{\hat{\boldsymbol{X}}_0^l, P_0^l\}$, $\forall i \in [1, N]$.
for $k = 0, 1, \ldots$ **do**
  Particle prediction: Draw the particles $\boldsymbol{X}_{k+1}^{n,i} \sim p(\boldsymbol{X}_{k+1}^n | \mathcal{X}_k^{n,i}, \mathcal{Y}_k)$ (3.41b).
  Kalman prediction: For each particles $\boldsymbol{X}_{k+1}^{n,i}$, sample the mean $\boldsymbol{X}_{k+1|k}^{l,i}$ and covariance $\boldsymbol{P}_{k+1|k}^{l,i}$ (3.38a)-(3.39b) of the density $p(\boldsymbol{X}_{k+1}^l | \mathcal{X}_{k+1}^n, \mathcal{Y}_k)$ (3.36b).
  Particle A2BC correction: Determine ϵ_k. Update the weights $w_{k+1}^i \propto w_k^i \, \mathcal{K}_{\epsilon_k}\left(\boldsymbol{Y}_{k+1} - h_{k+1}\left([\boldsymbol{X}_{k+1}^{n,i}, \boldsymbol{X}_{k+1|k}^{l,i}]^\top\right)\right)$ (3.41a).
  if a resampling criterion is satisfied, e.g., $\widehat{\mathrm{ESS}} < N_{\mathrm{th}}$, see (3.19) **then**
    Draw a new set of particles $\{\boldsymbol{X}_{k+1}^i = [\boldsymbol{X}_{k+1}^{n,i}, \boldsymbol{X}_{k+1|k}^{l,i}]^\top\}_{i\in[1,N]}$ using a resampling method, e.g., the multinomial resampling (Algorithm 3.2).
    Reset the weights to $\frac{1}{N}$.
  end if
  Kalman correction: Update the mean $\boldsymbol{X}_{k+1|k+1}^{l,i}$ and covariance $\boldsymbol{P}_{k+1|k+1}^{l,i}$ (3.37a)-(3.37d) of the density $p(\boldsymbol{X}_{k+1}^l | \mathcal{X}_{k+1}^n, \mathcal{Y}_{k+1})$ (3.36a).
  Compute the nonlinear state estimate $\hat{\boldsymbol{X}}_{k+1}^n$ (3.16) and its covariance $\hat{\boldsymbol{P}}_{k+1}^n$ (3.17).
  Compute the linear state estimate $\hat{\boldsymbol{X}}_{k+1}^l$ (3.42a) and its covariance $\hat{\boldsymbol{P}}_{k+1}^l$ (3.42b).
end for
Return the state estimate $\hat{\boldsymbol{X}}_{k+1} = [\hat{\boldsymbol{X}}_{k+1}^n, \hat{\boldsymbol{X}}_{k+1}^l]^\top$ and its covariance $\hat{\boldsymbol{P}}_{k+1} = \begin{pmatrix} \hat{\boldsymbol{P}}_{k+1}^n & 0_{n\times l} \\ 0_{l\times n} & \hat{\boldsymbol{P}}_{k+1}^l \end{pmatrix}$, $\forall k$.

BIBLIOGRAPHY

[1] Brian DO Anderson and John B Moore. *Optimal filtering*. Courier Corporation, 2012.

[2] Brian DO Anderson, Adrian N Bishop, Pierre Del Moral, and Camille Palmier. "Backward nonlinear smoothing diffusions". In: *Teor. Veroyatnost. i Primenen.* 66.2 (2021), pp. 305–326.

[3] Jeffrey L Anderson. "An ensemble adjustment Kalman filter for data assimilation". In: *Monthly weather review* 129.12 (2001), pp. 2884–2903.

[4] Jeffrey L Anderson. "An adaptive covariance inflation error correction algorithm for ensemble filters". In: *Tellus A: Dynamic meteorology and oceanography* 59.2 (2007), pp. 210–224.

[5] Jeffrey L Anderson and Stephen L Anderson. "A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts". In: *Monthly Weather Review* 127.12 (1999), pp. 2741–2758.

[6] Alan Bain and Dan Crisan. *Fundamentals of stochastic filtering*. Vol. 60. Springer Science & Business Media, 2008.

[7] Yaakov Bar-Shalom, X Rong Li, and Thiagalingam Kirubarajan. *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2004.

[8] Jacques Benichou and Mitchell H Gail. "A delta method for implicitly defined random variables". In: *The American Statistician* 43.1 (1989), pp. 41–44.

[9] Niclas Bergman. "Recursive Bayesian estimation: Navigation and tracking applications". PhD thesis. Linköping University, 1999.

[10] Yannick Bidel, Olivier Carraz, Renée Charriere, Malo Cadoret, Nassim Zahzam, and Alexandre Bresson. "Compact cold atom gravimeter for field applications". In: *Applied Physics Letters* 102.14 (2013), p. 144107.

[11] Craig H Bishop, Brian J Etherton, and Sharanya J Majumdar. "Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects". In: *Monthly weather review* 129.3 (2001), pp. 420–436.

[12] Miodrag Bolic. *Architectures for efficient implementation of particle filters*. State University of New York at Stony Brook, 2004.

[13] Ch J Bordé. "Atomic interferometry with internal state labelling". In: *Physics letters A* 140.1-2 (1989), pp. 10–12.

[14] Kenneth R Britting. "Inertial navigation systems analysis". In: (1971).

[15] Gerrit Burgers, Peter Jan van Leeuwen, and Geir Evensen. "Analysis scheme in the ensemble Kalman filter". In: *Monthly weather review* 126.6 (1998), pp. 1719–1724.

[16] Laurent Calvet and Veronika Czellar. "Accurate Methods for Approximate Bayesian Computation Filtering". In: (2012).

[17] George Casella and Christian P Robert. "Rao-Blackwellisation of sampling schemes". In: *Biometrika* 83.1 (1996), pp. 81–94.

[18] John F Clauser. "Ultra-high sensitivity accelerometers and gyroscopes using neutral atom matter-wave interferometry". In: *Physica B+ C* 151.1-2 (1988), pp. 262–272.

[19] Karim Dahia. "Nouvelles méthodes en filtrage particulaire-Application au recalage de navigation inertielle par mesures altimétriques". PhD thesis. Université Joseph-Fourier-Grenoble I, 2005.

[20] Pierre Del Moral. "Non Linear Filtering: Interacting Particle Solution". In: *Markov Processes and Related Fields* 2.4 (1996), pp. 555–580.

[21] Pierre Del Moral. "Measure-valued processes and interacting particle systems. Application to nonlinear filtering problems". In: *The Annals of Applied Probability* 8.2 (1998), pp. 438–495.

[22] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. "An adaptive sequential Monte Carlo method for approximate Bayesian computation". In: *Statistics and computing* 22.5 (2012), pp. 1009–1020.

[23] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. "On adaptive resampling strategies for sequential Monte Carlo methods". In: *Bernoulli* 18.1 (2012), pp. 252–278.

[24] Pierre Del Moral, Jean Jacod, and Philip Protter. "The Monte-Carlo method for filtering with discrete-time observations". In: *Probability Theory and Related Fields* 120.3 (2001), pp. 346–368.

[25] Pierre Del Moral and Spiridon Penev. *Stochastic Processes: From Applications to Theory*. Chapman and Hall/CRC, 2017.

[26] Luc Devroye. "Nonuniform random variate generation". In: *Handbooks in operations research and management science* 13 (2006), pp. 83–121.

[27] Arnaud Doucet, Simon Godsill, and Christophe Andrieu. "On sequential Monte Carlo sampling methods for Bayesian filtering". In: *Statistics and computing* 10.3 (2000), pp. 197–208.

[28] Geir Evensen. "Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics". In: *Journal of Geophysical Research: Oceans* 99.C5 (1994), pp. 10143–10162.

[29] Geir Evensen. *Data assimilation: the ensemble Kalman filter*. Springer Science & Business Media, 2009.

[30] Jl Galdos. "A Cramer-Rao bound for multidimensional discrete-time dynamical systems". In: *IEEE Transactions on Automatic Control* 25.1 (1980), pp. 117–119.

[31] J-L Gauvain and Chin-Hui Lee. "Maximum a posteriori estimation for multivariate Gaussian mixture observations of Markov chains". In: *IEEE transactions on speech and audio processing* 2.2 (1994), pp. 291–298.

[32] Christophe Giraud. "Introduction to high-dimensional statistics". In: *Monographs on Statistics and Applied Probability* 139 (2015), p. 139.

[33] Amadou Gning, Lyudmila Mihaylova, Fahed Abdallah, and Branko Ristic. "Particle filtering combined with interval methods for tracking applications". In: *Integrated Tracking, Classification, and Sensor Management* (2013), pp. 43–74.

[34] Neil J Gordon, David J Salmond, and Adrian FM Smith. "Novel approach to nonlinear/non-Gaussian Bayesian state estimation". In: *IEE proceedings F (radar and signal processing)*. Vol. 140. 2. IET. 1993, pp. 107–113.

[35] Peter J Green, Krzysztof Łatuszyński, Marcelo Pereyra, and Christian P Robert. "Bayesian computation: a summary of the current state, and samples backwards and forwards". In: *Statistics and Computing* 25.4 (2015), pp. 835–862.

[36] Ben Grocholsky. "Information-theoretic control of multiple sensor platforms". In: (2002).

[37] Fredrik Gustafsson, Fredrik Gunnarsson, Niclas Bergman, Urban Forssell, Jonas Jansson, Rickard Karlsson, and P-J Nordlund. "Particle filters for positioning, navigation, and tracking". In: *IEEE Transactions on signal processing* 50.2 (2002), pp. 425–437.

[38] Nicholas J Higham. "Computing the polar decomposition—with applications". In: *SIAM Journal on Scientific and Statistical Computing* 7.4 (1986), pp. 1160–1174.

[39] Nicholas J Higham. "Computing a nearest symmetric positive semidefinite matrix". In: *Linear algebra and its applications* 103 (1988), pp. 103–118.

[40] YC Ho and RCKA Lee. "A Bayesian approach to problems in stochastic estimation and control". In: *IEEE transactions on automatic control* 9.4 (1964), pp. 333–339.

[41] Peter L Houtekamer and Herschel L Mitchell. "A sequential ensemble Kalman filter for atmospheric data assimilation". In: *Monthly Weather Review* 129.1 (2001), pp. 123–137.

[42] *International Gravimetric Bureau*. http://bgi.obs-mip.fr/data-products/gravity-databases/marine-gravity-datas/. Visited on 2021-07-19.

[43] Ajay Jasra, Sumeetpal S Singh, James S Martin, and Emma McCoy. "Filtering via approximate Bayesian computation". In: *Statistics and Computing* 22.6 (2012), pp. 1223–1237.

[44] Rudolph Emil Kalman. "A new approach to linear filtering and prediction problems". In: (1960).

[45] Rickard Karlsson and Fredrik Gustafsson. "Bayesian surface and underwater navigation". In: *IEEE Transactions on Signal Processing* 54.11 (2006), pp. 4204–4213.

[46] Rickard Karlsson, Thomas Schön, and Fredrik Gustafsson. "Complexity analysis of the marginalized particle filter". In: *IEEE Transactions on Signal Processing* 53.11 (2005), pp. 4408–4411.

[47] Genshiro Kitagawa. "A Monte Carlo filtering and smoothing method for non-Gaussian nonlinear state space models". In: *Proceedings of the 2nd US-Japan joint seminar on statistical time series analysis*. 1993, pp. 110–131.

[48] Genshiro Kitagawa. "Monte Carlo filter and smoother for non-Gaussian nonlinear state space models". In: *Journal of computational and graphical statistics* 5.1 (1996), pp. 1–25.

[49] Augustine Kong, Jun S Liu, and Wing Hung Wong. "Sequential imputations and Bayesian missing data problems". In: *Journal of the American statistical association* 89.425 (1994), pp. 278–288.

[50] Deok-Jin Lee. "Nonlinear Bayesian filtering with applications to estimation and navigation". PhD thesis. Texas A&M University, 2005.

[51] Peter Jan van Leeuwen. "A variance-minimizing filter for large-scale applications". In: *Monthly Weather Review* 131.9 (2003), pp. 2071–2084.

[52] Erich L Lehmann and George Casella. *Theory of point estimation*. Springer Science & Business Media, 2006.

[53] Cornelius T Leondes, John B Peller, and Edwin B Stear. "Nonlinear smoothing theory". In: *IEEE Transactions on Systems Science and Cybernetics* 6.1 (1970), pp. 63–71.

[54] Tiancheng Li, Miodrag Bolic, and Petar M Djuric. "Resampling methods for particle filtering: classification, implementation, and strategies". In: *IEEE Signal processing magazine* 32.3 (2015), pp. 70–86.

[55] Jean-Michel Marin, Pierre Pudlo, Christian P Robert, and Robin J Ryder. "Approximate Bayesian computational methods". In: *Statistics and Computing* 22.6 (2012), pp. 1167–1180.

[56] Paul Marjoram, John Molitor, Vincent Plagnol, and Simon Tavaré. "Markov chain Monte Carlo without likelihoods". In: *Proceedings of the National Academy of Sciences* 100.26 (2003), pp. 15324–15328.

[57] Deborah K Meduna. *Terrain relative navigation for sensor-limited systems with application to underwater vehicles*. Stanford University, 2011.

[58] José Melo and Aníbal Matos. "Survey on advances on terrain based navigation for autonomous underwater vehicles". In: *Ocean Engineering* 139 (2017), pp. 250–264.

[59] Nicolas Merlinge. "State estimation and trajectory planning using box particle kernels". PhD thesis. Université Paris-Saclay; Coventry University, 2018.

[60] Christian Musso and Nadia Oudjane. "Regularisation schemes for branching particle systems as a numerical solving method of the nonlinear filtering problem". In: *Proceedings of the Irish Signals Conference, Dublin* (1998).

[61] Christian Musso, Nadia Oudjane, and Francois Le Gland. "Improving regularised particle filters". In: *Sequential Monte Carlo methods in practice*. Springer, 2001, pp. 247–271.

[62] Christian Musso, Bernard Sacleux, Alexandre Bresson, Jean-Michel Allard, Karim Dahia, Yannick Bidel, Nassim Zahzam, and Camille Palmier. "Terrain-aided navigation with an atomic gravimeter". In: *2019 22th International Conference on Information Fusion (FUSION)*. IEEE. 2019, pp. 1–8.

[63] *Naval Postgraduate School, Ocean Acoustics Laboratory*. https://www.oc.nps.edu/oc4270/index_files/Page417.htm. Visited on 2021-07-19.

[64] Per-Johan Nordlund. *Sequential Monte Carlo filters and integrated navigation*. Citeseer, 2002.

[65] Per-Johan Nordlund and Fredrik Gustafsson. "Marginalized particle filter for accurate and reliable terrain-aided navigation". In: *IEEE Transactions on Aerospace and electronic systems* 45.4 (2009), pp. 1385–1399.

[66] Gary W Oehlert. "A note on the delta method". In: *The American Statistician* 46.1 (1992), pp. 27–29.

[67] Nadia Oudjane. "Stabilité et approximations particulaires en filtrage non linéaire application au pistage". PhD thesis. Université Rennes 1, 2000.

[68] Nadia Oudjane and Christian Musso. "Progressive correction for regularized particle filters". In: *Proceedings of the Third International Conference on Information Fusion*. Vol. 2. IEEE. 2000, THB2–10.

[69] Nicolas Papadakis. "Assimilation de données images: application au suivi de courbes et de champs de vecteurs". PhD thesis. Université Rennes 1, 2007.

[70] D Pham, Karim Dahia, and Christian Musso. "A kalman-particle kernel filter and its application to terrain navigation". In: *Proceedings 6th ICIF* (2003), p. 24.

[71] Dinh Tuan Pham. "Stochastic methods for sequential data assimilation in strongly nonlinear systems". In: *Monthly weather review* 129.5 (2001), pp. 1194–1207.

[72] Michael K Pitt and Neil Shephard. "Filtering via simulation: Auxiliary particle filters". In: *Journal of the American statistical association* 94.446 (1999), pp. 590–599.

[73] Jonathan K Pritchard, Mark T Seielstad, Anna Perez-Lezaun, and Marcus W Feldman. "Population growth of human Y chromosomes: a study of Y chromosome microsatellites". In: *Molecular biology and evolution* 16.12 (1999), pp. 1791–1798.

[74] Paul Bui Quang, Christian Musso, and Francois Le Gland. "Particle filtering and the laplace method for target tracking". In: *IEEE Transactions on Aerospace and Electronic Systems* 52.1 (2016), pp. 350–366.

[75] Herbert E Rauch, F Tung, and Charlotte T Striebel. "Maximum likelihood estimates of linear dynamic systems". In: *AIAA journal* 3.8 (1965), pp. 1445–1450.

[76] Branko Ristic, Sanjeev Arulampalam, and Neil Gordon. *Beyond the Kalman filter: Particle filters for tracking applications*. Vol. 685. Artech house Boston, 2004.

[77] Pavel Sakov and Peter R Oke. "A deterministic formulation of the ensemble Kalman filter: an alternative to ensemble square root filters". In: *Tellus A: Dynamic Meteorology and Oceanography* 60.2 (2008), pp. 361–371.

[78] Simo Särkkä and Arno Solin. *Applied stochastic differential equations*. Vol. 10. Cambridge University Press, 2019.

[79] Thomas Schön. "On computational methods for nonlinear estimation". PhD thesis. Linköping University, 2003.

[80] Thomas Schön, Fredrik Gustafsson, and P-J Nordlund. "Marginalized particle filters for mixed linear/nonlinear state-space models". In: *IEEE Transactions on signal processing* 53.7 (2005), pp. 2279–2289.

[81] Bernard W Silverman. *Density estimation for statistics and data analysis*. Vol. 26. CRC press, 1986.

[82] Simon Tavaré, David J Balding, Robert C Griffiths, and Peter Donnelly. "Inferring coalescence times from DNA sequence data". In: *Genetics* 145.2 (1997), pp. 505–518.

[83] Francisco Curado Teixeira, João Quintas, Pramod Maurya, and António Pascoal. "Robust particle filter formulations with application to terrain-aided navigation". In: *International journal of Adaptive control and signal processing* 31.4 (2017), pp. 608–651.

[84] Petr Tichavsky, Carlos H Muravchik, and Arye Nehorai. "Posterior Cramér-Rao bounds for discrete-time nonlinear filtering". In: *IEEE Transactions on signal processing* 46.5 (1998), pp. 1386–1396.

[85] Michael K Tippett, Jeffrey L Anderson, Craig H Bishop, Thomas M Hamill, and Jeffrey S Whitaker. "Ensemble square root filters". In: *Monthly Weather Review* 131.7 (2003), pp. 1485–1490.

[86] Rudolph Van Der Merwe, Arnaud Doucet, Nando De Freitas, and Eric A Wan. "The unscented particle filter". In: *Tech. Rep. CUED/F-INFENG/TR 380, Cambridge University Engineering Department* (2000).

[87] Eric A Wan and Rudolph Van Der Merwe. "The unscented Kalman filter for nonlinear estimation". In: *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373)*. Ieee. 2000, pp. 153–158.